Platform Siloing and Soft Censorship: RT, PressTV, Facebook, and the New Geography of Internet Visibility

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

July 5, 2026

This paper examines a modern form of information control that does not resemble the older model of national firewalls, explicit blacklists, or direct state censorship. Using the circulation of RT and PressTV articles on a United States–based Facebook page as a test case, the paper analyzes how visibility can be restricted, localized, demoted, or contained without the user receiving a formal warning or experiencing an obvious block. The central argument is that contemporary internet censorship increasingly operates through layered jurisdictional, corporate, algorithmic, and social mechanisms rather than through simple denial of access. A post may remain visible to the author in Kansas while being invisible, deprioritized, or legally restricted for viewers in the European Union, United Kingdom, Canada, Australia, or other regions. This produces a condition best described as soft censorship or platform siloing: the speaker is not silenced, but the audience is quietly narrowed. The paper develops a framework for distinguishing legal prohibition, platform enforcement, geoblocking, shadow distribution, self-selection, and measurable public reach.

Keywords: internet censorship, platform siloing, soft censorship, RT, Russia Today, PressTV, Facebook moderation, Meta, geoblocking, sanctions, state media, algorithmic suppression, information warfare, digital sovereignty, social media regulation, EU media sanctions, online visibility, shadow banning, platform governance, foreign state media.

A collaboration with OpenAI GPT-5.5 Thinking. CC4.0.

https://orcid.org/0009-0004-4206-4253

Outline

1. Introduction: From Blocking to Containment

1.1. The Disappearance of Obvious Censorship

1.2. The User-Side Illusion of Visibility

1.3. RT and PressTV as Boundary Cases

1.4. Thesis: The Speaker May Be Visible to Himself but Not to the World

2. The U.S. Internet Baseline: Open Access Without a National Firewall

2.1. The Absence of a General Federal Internet Blacklist

2.2. Illegal Content, Domain Seizures, and Narrow Exceptions

2.3. Schools, Libraries, Workplaces, and Private Network Filters

2.4. Why U.S. Access Does Not Predict Foreign Visibility

3. RT and PressTV as Restricted State-Media Objects

3.1. RT After the Ukraine War: Broadcast Bans, Sanctions, and Platform Removal

3.2. PressTV After Iran Sanctions: Domain, Satellite, and Platform Disruption

3.3. The Difference Between a Website, a Broadcast Channel, and a Copied Article

3.4. Why “No URL” Does Not Necessarily Mean “No Restriction”

4. Legal Geography: How One Post Can Exist in One Country and Disappear in Another

4.1. European Union Media Restrictions and the Expansion of Distribution Liability

4.2. The United Kingdom and Canada: Broadcast Regulation Versus General Web Access

4.3. Australia, Sanctions, and Reputational Enforcement

4.4. Local Law Restriction as a Country-Specific Visibility Layer

5. Platform Governance: Facebook as a Private Border System

5.1. Meta’s State-Media Policies and Entity-Based Enforcement

5.2. Link Blocking Versus Text Recognition

5.3. Removal, Demotion, Warning Labels, and Reduced Distribution

5.4. Why Enforcement Can Be Real but Inconsistent

6. The Copied-Article Problem

6.1. Full Reposting Versus Quotation and Commentary

6.2. Copyright Risk Independent of Political Content

6.3. The Difference Between Archiving, Criticism, and Acting as a Mirror Site

6.4. How Platform Systems May Interpret Repeated Verbatim Reposting

7. Siloing as a Form of Soft Censorship

7.1. Defining the Silo

7.2. Geographic Siloing

7.3. Algorithmic Siloing

7.4. Social Siloing and Self-Selected Audiences

7.5. The Difference Between Being Blocked and Being Contained

8. The Psychology of the Invisible Wall

8.1. “I Have Had No Problem” as a Data Point

8.2. Why the Author’s View Is Not the Audience’s View

8.3. The Comfort of Apparent Normalcy

8.4. Containment Without Notification

9. Testing Visibility Without Creating False Signals

9.1. Why a VPN Is an Imperfect Test

9.2. Why a New Foreign Account Can Mislead the Experiment

9.3. Human Verification by Real Viewers in Target Countries

9.4. Logged-Out Checks, Direct Links, and Shareability Tests

9.5. Comparing Ordinary Posts, RT Posts, and PressTV Posts

10. Measurement Limits Inside the Facebook Dashboard

10.1. Missing Country-Level Data

10.2. Thresholds, Privacy Limits, and Page-Type Restrictions

10.3. Reach, Engagement, Shares, and Non-Follower Exposure

10.4. The Problem of Measuring What the Platform Does Not Report

11. A Taxonomy of Modern Internet Restriction

11.1. Hard Blocking

11.2. Legal Prohibition

11.3. Platform Removal

11.4. Geoblocking

11.5. Demotion and Distribution Reduction

11.6. Social Containment

11.7. User Unawareness as a Structural Feature

12. Conclusion: The Future of Censorship Is Partial Visibility

12.1. Why the Old Question “Is It Blocked?” Is No Longer Enough

12.2. The Rise of Jurisdictional and Algorithmic Fragmentation

12.3. Practical Implications for Researchers, Writers, and Citizens

12.4. Final Claim: The Post Exists, but the Public May Not

13. References

13.1. Legal and Regulatory Sources

13.2. Platform Policy Sources

13.3. Government Sanctions and Media Restriction Sources

13.4. Journalism and Transparency Reports

13.5. Academic Literature on Platform Governance and Soft Censorship

Art

1. Introduction: From Blocking to Containment

The older image of internet censorship was simple, visible, and often theatrical. A government blocked a website. A user entered a forbidden address. The browser failed. The citizen learned immediately that a border had been reached. This model still exists in many countries, but it no longer captures the most important form of information restriction in the networked democracies, regulated markets, and platform-governed public spheres of the twenty-first century. The new system is less dramatic. The post remains online. The author can see it. Friends may like it. A familiar audience may comment on it. No warning appears. No ban is issued. Yet the post may not travel, may not cross borders, may not appear in feeds, may not be shown to strangers, may not be shareable in certain jurisdictions, or may be placed into a low-visibility category that leaves the speaker unaware of the true size and shape of the audience.

This paper begins from a practical question: what happens when a United States–based Facebook user repeatedly copies and posts articles from RT or PressTV without posting direct URLs? The question is not merely about those two outlets. It is a window into a broader transformation in information control. RT and PressTV occupy a peculiar position in the modern media ecosystem. They are not ordinary foreign newspapers. They are state-linked media entities associated with geopolitical adversaries of the United States and its allies. They are subject to sanctions, platform restrictions, broadcast bans, domain seizures, app removals, labeling, demotion, and jurisdiction-specific rules. Yet their content can still appear in copied form, quoted form, summarized form, screenshot form, and reposted form. This creates a gray zone in which the content is neither fully banned nor fully free.

The central issue is visibility. A post does not need to be deleted to be controlled. A user does not need to be banned to be contained. A country does not need to operate a national firewall to shape what its citizens can see. A platform does not need to announce censorship to reduce distribution. In the contemporary system, restriction often occurs by narrowing reach rather than by destroying speech. The result is a new ambiguity. The author asks whether he is blocked; the more precise question is whether he is being silently limited, jurisdictionally partitioned, algorithmically deprioritized, or socially enclosed.

1.1. The Disappearance of Obvious Censorship

Obvious censorship is easy to recognize because it makes itself known. A book is banned. A broadcaster loses its license. A domain is seized. A website refuses to load. A user receives a notice that content is unavailable in a country because of local law. These forms remain important, and in some cases they are still decisive. But they are no longer the whole story. The most powerful platforms do not merely host speech; they rank, recommend, label, distribute, suppress, route, localize, and measure it. They decide whether a post appears at the top of a

feed, at the bottom of a feed, inside a small community, outside the network, in search results, in recommendations, or not at all.

This means censorship can disappear as an event while remaining present as a condition. The user does not experience a confrontation with authority. Instead, the user experiences continuity. The account remains active. The page remains open. Posts continue to publish. Engagement may continue at a low or familiar level. Nothing seems wrong. Yet the audience may have changed. The post may be traveling only to those who already interact with the page. It may not be offered to new viewers. It may not circulate internationally. It may be hidden in one country and visible in another. It may be permitted as speech but discouraged as distribution.

This disappearance of obvious censorship changes the psychology of the speaker. Under the old model, censorship produced certainty: the speaker knew he had been blocked. Under the new model, the speaker is left with uncertainty. Is the post unpopular? Is the audience small? Is the platform indifferent? Is the post being downranked? Is the content legally restricted in another country? Is the dashboard incomplete? Is the apparent absence of trouble evidence of freedom, or evidence that the post never escaped the silo?

1.2. The User-Side Illusion of Visibility

The most misleading feature of platform speech is that the author's view is not the public view. On a social platform, the author typically sees his own post as successful once it appears on his page. The post is there. It can be opened. It may show likes or comments. It can be copied. It may even be shareable. But the author's visibility is only one layer. The real question is not whether the post exists in the author's account. The real question is where the post is distributed, to whom it is shown, whether strangers can discover it, whether foreign viewers can access it, whether the post appears in feeds, and whether the platform treats it as normal content.

This illusion is especially strong for users who post daily and develop a stable audience. If the same group of people repeatedly sees and interacts with the material, the author may conclude that nothing is being restricted. But repeated engagement from a stable circle does not prove broad public visibility. It may prove the opposite. It may show that the content has been enclosed inside an audience that the platform already understands as receptive, tolerant, or politically compatible. The post is not deleted; it is domesticated.

The dashboard often deepens the uncertainty. A platform may show impressions, reach, reactions, and comments without revealing the most important missing data. It may not disclose country-level visibility. It may not show whether the post was suppressed in feeds. It may not identify whether non-followers were excluded. It may not reveal whether a particular

topic, phrase, source, or copied article triggered reduced distribution. The user receives enough data to believe he is informed, but not enough data to audit the distribution system.

The result is a user-side illusion: because the post is visible to the author, the author assumes it is visible to the world. But platform visibility is not binary. It is scalar, relational, and jurisdictional. A post may be visible to the author, visible to followers, visible to domestic viewers, invisible to foreign viewers, absent from recommendations, blocked from sharing, unavailable to logged-out users, or hidden from search. The object called "the post" is therefore not one object. It is many platform states presented behind a single interface.

1.3. RT and PressTV as Boundary Cases

RT and PressTV are useful boundary cases because they reveal the difference between speech, source, distribution, and jurisdiction. In the United States, a private person can generally read, discuss, quote, criticize, or analyze foreign state media. The American internet does not operate through a broad national firewall that blocks ordinary access to every disfavored foreign news source. Yet RT and PressTV are not treated as ordinary sources by many governments and platforms. RT has faced broadcast bans, platform removals, and legal restrictions in several Western jurisdictions following Russia's invasion of Ukraine. PressTV has faced sanctions, domain disruptions, satellite removals, and platform restrictions linked to the Iranian state and its information apparatus.

The copied-article problem complicates enforcement. If a user posts a direct RT link, the platform can identify the domain. If the user copies the article text without a URL, the platform must treat the content differently. It may ignore it. It may detect the text. It may classify it by source language, distinctive phrasing, metadata, user behavior, reports, or repeated pattern. It may treat the post as commentary if surrounded by analysis. It may treat it as republication if copied in full. It may treat occasional quotation differently from daily mirroring. The absence of a URL does not necessarily remove the association with the source. It merely shifts detection from domain-level control to content-level and behavior-level interpretation.

These outlets therefore sit at the edge of several regimes at once. They are media sources, state instruments, copyrighted publishers, geopolitical signals, sanctions objects, and platform-policy triggers. A Facebook post containing their text may be treated differently depending on location, format, frequency, context, and audience. A short excerpt used in criticism is not the same as a full article copied daily. A post visible in Kansas may not be visible in Berlin. A PressTV article may trigger a different system than an RT article. A human reader may see journalism; an automated system may see state-media amplification.

Boundary cases matter because they expose the architecture of the system. Ordinary content may pass through the platform unnoticed. Prohibited content may be removed. But politically

sensitive state-media content often falls into the intermediate zone where visibility is managed, not simply granted or denied. That intermediate zone is where the modern politics of platform speech increasingly resides.

1.4. Thesis: The Speaker May Be Visible to Himself but Not to the World

The thesis of this paper is that contemporary internet restriction increasingly operates through partial visibility rather than total erasure. The decisive question is no longer simply whether a source is blocked. The better question is how, where, by whom, for whom, and at what layer it is restricted. A post may survive at the level of authorship while failing at the level of circulation. It may remain on the page while disappearing from foreign audiences. It may be legally tolerated in one country and suppressed in another. It may be technically accessible but algorithmically buried. It may be formally public yet practically private.

This condition can be called platform siloing. A silo is not a prison in the ordinary sense. The speaker can still speak. The page can still publish. The archive can still grow. But the audience is bounded by invisible walls. Those walls may be legal, algorithmic, social, commercial, or geopolitical. They may be raised by the platform, by the state, by user behavior, by reporting systems, by sanctions compliance, by copyright systems, or by the recommender algorithm's assessment of risk. Their power lies in the fact that they do not need to announce themselves.

Soft censorship is therefore not the abolition of speech. It is the management of reach. It leaves the symbolic surface of freedom intact while altering the practical conditions of communication. The user sees continuity; the public may see absence. The speaker experiences no punishment; the message fails to travel. The platform avoids the appearance of banning; the distribution system produces containment.

The practical implication is serious. Researchers, writers, journalists, and citizens can no longer infer freedom from publication alone. To publish is not necessarily to reach. To be visible to oneself is not necessarily to be visible to others. To avoid removal is not necessarily to avoid restriction. The post exists, but its public may not. This paper develops that claim through the case of RT, PressTV, Facebook, and the fragmented geography of modern internet visibility.

2. The U.S. Internet Baseline: Open Access Without a National Firewall

The United States is often criticized for surveillance, corporate platform power, opaque moderation, sanctions pressure, and informal influence over digital infrastructure. Those criticisms are not trivial. They describe real features of the American information environment. Yet the United States does not normally operate the internet through the classic architecture of a national firewall. There is no ordinary federal system that blocks citizens from reading foreign

news, foreign state media, dissident political writing, controversial religious material, rival economic analysis, anti-government commentary, or unpopular ideological arguments. A user in Kansas can generally access Russian, Iranian, Chinese, British, French, Israeli, Arab, Indian, African, Latin American, and European media sources directly from a normal browser unless some other layer intervenes.

This baseline matters because it distinguishes access from distribution. A U.S.-based Facebook user may be able to read RT, PressTV, or other foreign state-linked outlets, copy article text, and post commentary without encountering a national block. But that does not mean the same material will circulate freely once placed inside a global platform. The American baseline is open at the network level but heavily mediated at the platform level. The user's browser may reach a website, but the user's Facebook post passes through a different governance system: private policy, automated detection, sanctions compliance, copyright enforcement, content ranking, local-law restrictions, and jurisdiction-specific display decisions.

The old question, "Is this website blocked in America?" is therefore too narrow. For the United States, the more useful answer is usually: no, not in the broad national-firewall sense. But the post may still be affected by criminal law, civil liability, sanctions law, copyright law, private platform rules, institutional filters, employer restrictions, school and library filters, app-store policies, payment restrictions, advertising rules, search ranking, or foreign visibility controls. The American system does not primarily prevent adults from reaching ordinary foreign information sources. It instead creates a layered environment in which certain kinds of content become risky, restricted, demonetized, delisted, deprioritized, localized, or unavailable in specific institutional and platform contexts.

This distinction is essential for understanding RT and PressTV reposting. In the United States, the author may experience the internet as open because direct access is often available. He may see the post publish successfully. He may receive likes and comments. He may not receive a warning. But this U.S. experience is not evidence that the post is globally visible. The absence of a national firewall does not eliminate platform siloing. It simply means the first gate is open.

2.1. The Absence of a General Federal Internet Blacklist

The United States has no ordinary public list of political websites that internet service providers must block for all citizens. There is no American equivalent of a centralized national censorship firewall that routinely prevents adults from reaching disfavored news organizations, foreign propaganda outlets, opposition parties, religious movements, ideological journals, or hostile foreign broadcasters. This is partly constitutional, partly cultural, partly technological, and partly historical. American law strongly protects speech, especially political speech, even when that speech is offensive, false, foreign, propagandistic, or hostile to prevailing opinion. The state may

regulate conduct, crime, fraud, espionage, threats, and specific unlawful transactions, but a general ban on reading or discussing foreign state media would face serious constitutional and practical obstacles.

This does not mean the United States is neutral about foreign influence. It is not. The U.S. government can designate entities, impose sanctions, seize domains, prosecute crimes, regulate foreign agents, restrict transactions, remove foreign propaganda from government devices, or pressure platforms through hearings, public criticism, law-enforcement warnings, intelligence assessments, or regulatory threats. These tools can be powerful. But they usually do not take the simple form of preventing every private citizen from typing a foreign news URL into a browser.

The distinction between access and endorsement is crucial. A citizen may be able to access a Russian or Iranian outlet without the U.S. government endorsing that outlet. A platform may allow a user to discuss a state-media article while labeling, demoting, or restricting the source. A court may tolerate publication while allowing sanctions on financial services or broadcast distribution. The legal system can permit reading while restricting commercial collaboration, broadcast licensing, foreign-agent concealment, or material support in specific contexts.

This produces the paradox of the American information environment. The user can often reach almost anything, but the surrounding ecosystem may attach consequences to repetition, amplification, monetization, coordination, or institutional distribution. A person reading RT is one thing. A broadcaster carrying RT is another. A platform algorithm recommending RT-derived content is another. A paid campaign amplifying RT narratives is another. A user copying entire RT articles daily as a substitute distribution path may be interpreted differently from a user quoting paragraphs for criticism. The absence of a national blacklist does not mean all forms of distribution are treated alike.

For this paper, the key point is that the U.S. baseline gives the author a sense of normal access. From inside the United States, the post may appear lawful, ordinary, and technically successful. But the same post exists inside a platform whose audience is transnational. Once a U.S. post becomes visible to viewers in the European Union, the United Kingdom, Canada, Australia, or other jurisdictions, it may encounter rules that do not exist in the same way at the author's physical location. The post is born in one legal environment but travels through many.

2.2. Illegal Content, Domain Seizures, and Narrow Exceptions

Although the United States lacks a general political firewall, it does not allow all internet content equally. Certain categories of material are criminal, actionable, or subject to seizure. The clearest example is child sexual abuse material. Other areas include criminal marketplaces, fraud schemes, counterfeit goods, piracy operations, malware infrastructure, terror-related

material in certain contexts, and domains used for sanctions evasion or unlawful transactions. In these cases, the state may seize domain names, prosecute operators, require removal, or coordinate with service providers.

These actions are different from general censorship in several ways. First, they usually target material or conduct already unlawful under existing criminal or civil law. Second, they are often case-specific rather than universal. Third, they frequently focus on operators, infrastructure, money flows, or distribution networks rather than the mere act of reading controversial political information. Fourth, they may produce visible signs, such as a federal seizure banner, rather than invisible distribution reduction. A seized domain announces state action. A demoted post does not.

Domain seizure is also not the same as total internet erasure. When a domain is seized, the content may move, reappear, mirror, or circulate elsewhere. The seizure affects a named address or infrastructure point. It does not necessarily remove all copies of the underlying text from the internet. This distinction matters for RT and PressTV because state-linked media content can move across formats. A website can be blocked, a YouTube channel removed, a satellite feed discontinued, an app delisted, a domain seized, or an account banned; yet the text of an article can still be copied, quoted, screenshotted, summarized, or reposted by private users.

The United States also treats informational materials differently from services or transactions in some sanctions contexts. The circulation of books, films, news articles, photographs, and other completed informational materials has historically received special protection. That does not mean every interaction with a sanctioned media entity is permitted. Paying, coordinating, advertising, hosting, consulting, producing commissioned content, or providing services can raise very different issues from reading or discussing already-published material. But the distinction helps explain why a U.S.-based user may not experience direct trouble for posting foreign-media text, especially if no money, coordination, or formal service relationship exists.

Still, narrow exceptions can be misunderstood. A user may infer that because he has not been blocked, the material is unrestricted everywhere. That conclusion does not follow. U.S. criminal law, sanctions law, copyright law, platform policy, and foreign local-law restrictions operate through different thresholds. A post can be lawful to read, risky to copy in full, contrary to platform policy, unavailable in some countries, and algorithmically demoted at the same time. Modern restriction is not one switch; it is a stack of switches.

2.3. Schools, Libraries, Workplaces, and Private Network Filters

The United States does contain many local and institutional filters. Schools and libraries may use internet filtering systems, especially when public funding conditions apply. Employers may block

social media, adult content, gambling sites, file-sharing services, foreign media, political websites, or cybersecurity risks. Universities may restrict malware, piracy, phishing, or high-risk domains. Public Wi-Fi networks may filter categories of sites for liability, bandwidth, institutional policy, or child-protection reasons. Parents may install household filters. Internet service providers may offer optional security filters. None of these systems amounts to a national firewall, but they shape the internet that many users actually experience.

These filters are often category-based rather than source-specific. A school may block adult material, weapons content, gaming, social platforms, anonymous proxies, or file-sharing. A workplace may block streaming video, political sites, foreign domains, or anything classified as productivity risk. A library may apply child-protection filtering on public terminals. A security gateway may block domains associated with malware or phishing. In these contexts, a user may conclude that a site is “blocked in America” when in fact it is blocked only on that network.

Institutional filtering introduces another lesson for platform siloing: internet restriction is often local, contextual, and invisible to outsiders. A website may be accessible from a home connection but inaccessible at a school. A Facebook post may be visible in the United States but unavailable in Germany. A copied RT article may be readable by the author but not recommended to non-followers. A PressTV post may survive on a page but be hidden from search. The user sees one path through the network and mistakes it for the whole network.

Private filters are also less accountable than government censorship. A school, company, or platform may not provide detailed explanations for every restriction. The user may see a vague message, a failed load, a missing post, or no indication at all. This resembles platform demotion more than classic censorship. The restriction may not announce its reason. It may be justified as safety, compliance, quality, spam reduction, authenticity enforcement, intellectual property protection, foreign-interference prevention, or user experience optimization.

For the American Facebook user, the lesson is that the U.S. internet is open but not homogeneous. There is no single American internet experience. There are home networks, mobile networks, workplace networks, school networks, library networks, platform networks, app-store networks, and payment networks. Each has its own governance layer. A user posting from home may be operating in the least restricted part of the stack while his potential audience encounters more restrictive layers elsewhere.

2.4. Why U.S. Access Does Not Predict Foreign Visibility

The most important error is to treat U.S. access as evidence of global visibility. A post created in the United States can enter foreign legal environments as soon as it becomes visible to foreign users. Facebook is not merely a U.S. bulletin board. It is a global distribution system that must respond to many jurisdictions, each with different rules governing hate speech, defamation,

state media, sanctions, privacy, copyright, election interference, terrorist content, extremist material, children's safety, and national security. The platform therefore does not need to apply one rule everywhere. It can localize restrictions.

This means a U.S. author can be fully able to publish while a viewer in another country is unable to see, share, search, or interact with the post. The author may see no warning because the content is not restricted for him. The post may remain in the archive because it does not violate global platform rules. But it may be unavailable in one jurisdiction because of local law, restricted in another because of sanctions risk, demoted generally because of state-media association, or shown only to followers because of political-content ranking decisions.

RT illustrates this problem sharply. A U.S. user may be able to find RT material, copy text, and discuss it. But RT has been treated in parts of Europe not merely as a controversial outlet, but as a restricted instrument of Russian state information warfare. That legal and political classification changes the distribution environment. If a platform believes a post is functionally redistributing RT content to European users, it may choose to restrict visibility in those countries even if the same post remains visible to the author in the United States. The user experiences publication. The foreign viewer experiences absence.

PressTV presents a related but distinct case. Its restrictions often arise through sanctions, platform removals, domain actions, satellite disruptions, and reputational enforcement rather than the same pattern of broad EU wartime media prohibition associated with RT. Yet the underlying principle is similar. A source may be accessible in one country, sanctioned in another, removed from a platform in another, and treated as state-linked or risky by automated systems everywhere. The copied text carries source identity even when the URL is absent.

The central point is that American openness is only the first condition of speech. It answers the question, "Can the author reach and post the material?" It does not answer the questions, "Who receives it?", "Where is it suppressed?", "Does the platform recommend it?", "Can non-followers discover it?", "Can foreign viewers share it?", or "Does the system treat it as ordinary speech or risky amplification?" For modern information circulation, those later questions are often more important than the first.

The U.S. internet baseline therefore creates a false sense of universality. Because the author can see the post, he assumes the post exists for everyone. Because no federal firewall blocks him, he assumes there is no restriction. Because no warning appears, he assumes there is no enforcement. But platform governance and international law have made visibility conditional. The post may be real in the United States, partial in Canada, legally sensitive in Europe, algorithmically disfavored across the platform, and socially contained everywhere. The future of

censorship is not necessarily the blocked page. It is the page that opens for the author and fails to arrive for the public.

3. RT and PressTV as Restricted State-Media Objects

RT and PressTV are not ordinary foreign media sources in the platform-governance environment. They occupy a special category because they are associated with states that Western governments often treat as adversarial information powers. This does not mean that every article they publish is false, nor that every reader is acting as an agent of a foreign government. It means that the source identity itself changes the regulatory and platform context. A copied article from RT or PressTV is not merely text. It carries the institutional identity of the outlet, the geopolitical status of the state behind it, the history of sanctions or restrictions attached to that outlet, and the platform's own risk model for foreign influence operations.

This distinction is essential. A private American user may think he is posting an article, while the platform may classify the same object as state-media amplification. A reader may interpret the post as evidence, commentary, dissent, or alternative news, while a government regulator may interpret the source as an instrument of information warfare. A copyright system may interpret the same post as unauthorized republication. A sanctions-compliance system may treat the media entity as legally sensitive. An algorithm may not understand the user's intent at all; it may simply recognize copied text, repeated source patterns, or audience behavior associated with politically charged content.

RT and PressTV therefore show how modern restriction attaches not only to words but to institutional provenance. The words may be reproduced without a hyperlink. The URL may be absent. The post may contain no explicit instruction to visit the source. Yet the content may still be recognizable as originating from a restricted or sensitive state-media outlet. This makes the copied-article case more complicated than a simple link-blocking problem. The issue is not merely whether Facebook allows a domain. The issue is whether the platform, the jurisdiction, or the public interprets the copied text as redistribution of a restricted state-media object.

A free-speech analysis must keep these distinctions clear. Reading RT is not the same as working for RT. Quoting PressTV is not the same as providing services to PressTV. Criticizing an article is not the same as mirroring an outlet. Posting a paragraph is not the same as copying entire articles every day. A scholar analyzing propaganda is not the same as an account built to launder state-media material into new distribution channels. Yet automated systems do not always distinguish these cases well. The practical risk arises precisely in that gray zone where ordinary

political speech, quotation, criticism, republication, sanctions compliance, copyright law, and platform foreign-interference policy overlap.

3.1. RT After the Ukraine War: Broadcast Bans, Sanctions, and Platform Removal

RT became one of the clearest cases of Western state-media restriction after Russia's full-scale invasion of Ukraine in February 2022. Before that war, RT had already been controversial. Critics described it as a Russian state propaganda outlet designed to undermine Western confidence, amplify internal divisions, and present Russian geopolitical narratives in the style of international news. Supporters or sympathetic readers often described it differently: as a counterweight to Western media consensus, a source for neglected perspectives, or a platform willing to question U.S., NATO, and European policy. Both perceptions mattered because RT's power came partly from occupying the ambiguous space between journalism, commentary, propaganda, and geopolitical messaging.

After the invasion of Ukraine, Western governments and platforms increasingly stopped treating RT as merely a controversial broadcaster. The outlet became part of a wartime information-control regime. In Europe, RT faced broad restrictions aimed at preventing Russian state media from broadcasting or distributing content into the European public sphere. The legal theory was not that every RT sentence was individually unlawful. The theory was that the outlet functioned as an instrument of state-directed information warfare during an armed conflict. That shifted the regulatory focus from evaluating individual claims to restricting institutional distribution.

This distinction is crucial for copied text. If a government or platform restricts RT as an entity, the restriction may not depend entirely on whether a user posts a clickable RT URL. The concern is the circulation of RT content. A platform can remove official RT accounts, block RT domains, demote RT links, label Russian state-controlled media, prevent advertising, remove apps, or reduce recommendation of content that functions as RT redistribution. Once the entity is restricted, the platform's enforcement problem becomes broader than domain blocking. It must decide whether derivative circulation also matters.

The platform environment also changed. Major social-media companies removed or restricted RT accounts, labeled Russian state-controlled media, reduced distribution of related content, or blocked access in certain regions. These actions were partly legal, partly reputational, partly political, and partly business decisions. A platform operating globally had to avoid violating EU restrictions, avoid appearing to amplify Russian wartime messaging, maintain relations with governments, and protect its public image. The result was not a single uniform ban visible to every user, but a layered system of removals, regional blocks, downranking, and source-based sensitivity.

For the U.S.-based Facebook user, this creates a paradox. The user may still be able to locate RT material on the open web, copy it, paste it, and see it appear on his page. That experience suggests freedom. Yet the same post may be treated differently once it attempts to circulate to users in Europe or other jurisdictions where RT's distribution has been legally or politically restricted. The user's domestic experience therefore cannot reveal the post's international distribution state. RT's post-Ukraine status makes it a prime example of a source that may remain readable in one place while becoming distribution-restricted elsewhere.

RT's case also illustrates why the old term "blocked" is insufficient. RT can be blocked from broadcasting, removed from app stores, restricted on social platforms, unavailable on certain satellite systems, demoted in feeds, inaccessible in some jurisdictions, and still visible through mirrors, archives, copied text, screenshots, search snippets, commentary, and reposts. A user asking whether RT is blocked is therefore asking a question with many answers. It may be blocked as a broadcaster, restricted as an app, demoted as a link, banned as an account, accessible as a website, visible as copied text, and contained as a Facebook post.

3.2. PressTV After Iran Sanctions: Domain, Satellite, and Platform Disruption

PressTV occupies a related but different position. It is associated with the Iranian state and has long been treated by critics as part of Iran's international messaging system. Its restrictions have often emerged through sanctions, broadcast-license decisions, satellite distribution conflicts, domain seizures, platform removals, and accusations involving forced confessions, disinformation, or state propaganda. Unlike RT after the Ukraine invasion, PressTV's restriction pattern has not always taken the same form of broad European wartime media prohibition. Its status is instead shaped by the longer conflict between Iran and Western governments, especially over sanctions, nuclear policy, regional conflict, human rights, and information operations.

This difference matters because RT and PressTV should not be collapsed into one category. Both are state-linked media. Both have faced Western restrictions. Both can be treated as propaganda outlets by adversarial governments. But their legal histories, enforcement mechanisms, and geopolitical triggers differ. RT's recent restriction environment is strongly tied to Russia's war in Ukraine and the European response to Russian state media. PressTV's restriction environment is more closely tied to sanctions regimes, Iranian state control, human-rights concerns, and platform policies toward Iranian state media and affiliated networks.

PressTV's distribution has been disrupted in several ways over time. It has lost broadcast access in some jurisdictions. Its domains or affiliated Iranian media domains have been targeted in enforcement actions. Its satellite carriage has been affected. Its social-media presence has been restricted or removed at various points. These measures show that control can occur through

infrastructure rather than direct reader prohibition. A channel can lose satellite distribution without every article becoming illegal to read. A domain can be seized while copied articles remain available elsewhere. A YouTube account can disappear while the outlet's text circulates on other platforms. Sanctions can restrict transactions with an entity without creating a simple rule that no citizen may discuss its content.

For a Facebook user copying PressTV articles, the relevant risk is therefore more subtle than a straightforward web block. The post may not be automatically treated exactly like RT content. It may not trigger the same European wartime distribution prohibition. But it still carries a sensitive source identity. It may be recognized as Iranian state media. It may be affected by platform rules against coordinated inauthentic behavior, state-controlled media manipulation, sanctions-sensitive entities, or content associated with banned accounts. It may also raise copyright issues if reposted in full. The user may experience no visible problem because the enforcement systems are inconsistent, because the post remains within a small audience, because no one reports it, or because the platform does not classify the copied text with high confidence.

PressTV also reveals how sanctions and speech interact uneasily. Sanctions are often aimed at money, services, property, institutional relationships, and material support rather than private reading. A user who reads or criticizes an article is not obviously doing the same thing as a company providing paid distribution services to a sanctioned media entity. But platforms operate conservatively because they face legal, financial, and reputational risk. They may restrict or remove official accounts, block monetization, prevent advertising, or avoid distributing material that appears to support a sanctioned entity. The result is again partial visibility rather than a single clear prohibition.

PressTV's case is therefore best understood as a distributed restriction environment. The outlet may be limited by sanctions, satellite access, platform bans, domain actions, search visibility, payment restrictions, and reputational treatment. But the copied text of a PressTV article can still move through private accounts, screenshots, email lists, archives, and commentary posts. That movement does not erase the source identity. It simply transfers the enforcement problem from the official outlet to the secondary distributor.

3.3. The Difference Between a Website, a Broadcast Channel, and a Copied Article

Modern information control depends on the form in which content appears. The same article can exist as a webpage, a broadcast segment, a video clip, a social-media post, a screenshot, a PDF, an email, a quotation, a summary, or copied plain text. Each form passes through different legal and technical systems. Treating them as identical creates confusion.

A website is a network location. It has a domain, hosting infrastructure, search visibility, server logs, links, metadata, and jurisdictional exposure. A government or platform can block a domain, seize it, delist it, refuse to resolve it, label it, or warn users before they visit it. A website is relatively easy to classify because it identifies its source through domain structure. If a user posts a URL to RT or PressTV, the platform can recognize that domain immediately. The link can be blocked, labeled, deprioritized, or removed without reading the full article.

A broadcast channel is different. It depends on licenses, satellite carriage, cable distribution, app distribution, and platform availability. Governments have historically regulated broadcasting more heavily than private reading because broadcasting uses scarce spectrum, licensed infrastructure, public-interest obligations, or regulated distributors. A country may revoke a broadcaster's license or remove it from authorized distribution without criminalizing the private act of reading the broadcaster's text on the internet. This is why a statement such as "RT is banned" requires clarification. Banned as what? A television channel? A licensed broadcaster? A satellite service? A social-media account? A website? A source in search results? A link on a platform? A copied article?

A copied article is different again. Once the text is pasted into Facebook without a URL, the domain-level signal disappears from the surface. But the source may remain detectable through exact text matching, distinctive phrasing, embedded names, headlines, formatting, metadata, user captions, repeated behavior, or audience reports. The copied article also changes the legal issue. It may become copyright republication. It may look like mirroring. It may be treated as user-generated content rather than direct outlet distribution. It may be harder to classify automatically. It may pass unnoticed. It may be removed only after complaint. It may be tolerated when excerpted but disfavored when copied in full.

This distinction helps explain why a user can post daily without trouble. The platform may be more aggressive against official accounts and direct links than against copied text. It may not have enough confidence to classify every pasted article. It may not want to remove political speech without context. It may rely on reports. It may enforce only at scale. It may not prioritize a small page. Or it may simply demote the material silently rather than delete it. In all these cases, the user experiences normal posting while the distribution environment remains uncertain.

The difference between forms also matters for legitimacy. A person who quotes two paragraphs from RT to criticize Russian framing is engaged in analysis. A person who copies full RT articles every day without commentary may appear to be replacing the restricted source's distribution channel. A person who archives foreign propaganda for research may have a different purpose from a person who uses copied articles to bypass platform bans. Platforms cannot easily infer

intent, so they often rely on patterns. Frequency, completeness, attribution, commentary, and audience behavior all shape the interpretation.

Thus the question is not merely whether RT or PressTV is blocked. The question is what object is being distributed. A domain can be blocked while text survives. A broadcast channel can be banned while an article remains readable. A platform account can be removed while a private user reposts its content. A copied article can escape link filters while entering copyright and state-media amplification risk. Modern censorship is form-sensitive.

3.4. Why "No URL" Does Not Necessarily Mean "No Restriction"

Removing the URL changes the enforcement surface but does not necessarily remove the restriction. A URL is the easiest signal for a platform to detect, but it is not the only signal. Platforms can recognize copied text, repeated headlines, known source phrases, images, screenshots, account behavior, sharing patterns, user reports, and coordinated amplification. They can also infer risk from context. If a page repeatedly posts full articles from the same restricted outlet, the absence of a URL may not prevent classification. It may instead make the behavior appear more like circumvention.

The "no URL" assumption rests on an older understanding of internet control. In that model, blocking means preventing access to a specific address. If the address is absent, the block fails. But platform governance is not limited to address blocking. It includes content fingerprinting, natural-language processing, entity recognition, copyright matching, spam detection, political-content ranking, community reporting, and account-level behavior analysis. A copied article is therefore not invisible. It is merely less explicit.

There is also a legal and policy difference between quoting and laundering. If a user writes, "RT claims X; here is why that claim matters," and includes a short excerpt, the post presents itself as commentary. If a user copies the entire article every day, the post may function as an alternative distribution node. The more complete, frequent, and systematic the reposting becomes, the more likely it is to be interpreted as mirroring rather than discussion. The platform does not need to prove formal coordination with RT or PressTV to reduce distribution. It may simply decide that the pattern resembles state-media amplification, spam, copyright infringement, or low-quality repeated content.

The absence of a URL can also reduce transparency for readers. A direct link lets readers identify the source, evaluate the outlet, inspect the original context, and see the institutional provenance. Copied text without a link may detach content from its origin. That detachment can be innocent, especially when a user names the source clearly and comments on it. But it can also resemble a laundering process in which state-media narratives circulate without the visible

state-media label. Platforms concerned with foreign influence may be sensitive to that possibility.

At the same time, restriction is not guaranteed. A pasted RT or PressTV article may remain visible indefinitely. The platform may not detect it. It may detect it but ignore it. It may classify it as user commentary. It may demote it rather than remove it. It may only restrict it in certain countries. It may wait for a copyright complaint. It may treat PressTV differently from RT. It may act only when the page reaches a larger audience. Enforcement inconsistency is not evidence that no policy exists. It is a normal feature of large-scale platform governance.

The central conclusion is that source identity can survive format change. RT content does not cease to be RT-derived because the link is removed. PressTV text does not become ordinary user speech simply because it is pasted into Facebook. The post becomes a hybrid object: user speech carrying state-media content. That hybrid status is exactly what makes the modern visibility problem difficult. The author may be speaking in his own name, but the material may still trigger systems designed for restricted sources.

This is why the user's year-long absence of trouble is meaningful but not decisive. It shows that the account has not been obviously banned or the posts routinely removed. It does not show that the posts are broadly distributed, visible in every country, recommended to non-followers, protected from future enforcement, or interpreted neutrally by the platform. "No URL" may reduce friction, but it does not guarantee reach. The content may still be contained by the invisible architecture of platform governance.

4. Legal Geography: How One Post Can Exist in One Country and Disappear in Another

The internet appears global because the interface is global. A Facebook post created in Kansas seems to occupy a single public space. The author publishes once, sees one object, and assumes that object has entered the world. But the platform does not actually operate as one jurisdiction. It operates as a stack of national, regional, corporate, and technical jurisdictions layered behind a common interface. The same post may be visible in the United States, restricted in one European Union member state, unavailable in another country, demoted globally, removed from search, or visible only to logged-in users. The author sees a single post; the platform manages many versions of that post.

This is the problem of legal geography. A digital object does not remain inside the legal environment where it was created. Once it is distributed, it crosses borders without physically moving. A Facebook post written under the assumptions of American free-speech culture can immediately become subject to European sanctions policy, British broadcast norms, Canadian distribution rules, Australian sanctions law, copyright claims, platform moderation policy, and

local-law demands from governments. The post has no passport, but the platform must treat it as if it enters many jurisdictions at once.

RT and PressTV reveal this legal geography because they carry the identity of foreign state-linked media. A user may be in the United States, but the audience may be in Germany, France, Poland, the United Kingdom, Canada, Australia, or elsewhere. Those countries do not necessarily treat Russian and Iranian state media in the same way the United States treats a private citizen's reading habits. Some jurisdictions focus on broadcast distribution. Others focus on sanctions. Others focus on platform availability, satellite transmission, state propaganda, or harmful foreign influence. The result is a fragmented legal environment in which access, publication, and visibility diverge.

A post can therefore exist in one country and disappear in another without contradiction. The disappearance may not be visible to the author. If the platform restricts access only for viewers in a particular jurisdiction, the author may receive no warning. If the platform demotes the post rather than removes it, there may be no visible sign at all. If the platform decides the post is lawful in the United States but risky in the European Union, it may preserve the post domestically while limiting its reach abroad. This is not the old censorship model of a single ban. It is a distributed compliance model.

Legal geography also complicates personal responsibility. The author may not intend to distribute restricted content into a country. He may simply post to his own page. But platforms distribute by default. A public page is not a local bulletin board; it is a global publication mechanism. If foreign users can view the page, the post may be treated as available in their jurisdiction. This creates a practical tension between the author's intent and the platform's legal exposure. The author asks, "Can I post this?" The platform asks, "Where might this be seen, and what happens to us if it is seen there?"

The answer is rarely absolute. A platform may decide that a post is acceptable globally. It may remove it everywhere. It may restrict it only in certain countries. It may reduce distribution without deleting it. It may leave it alone until a government complaint, user report, copyright notice, or sanctions review occurs. It may treat RT differently from PressTV. It may treat direct links differently from copied text. It may treat full articles differently from excerpts with commentary. The legal geography of modern platforms is not only territorial; it is also contextual.

4.1. European Union Media Restrictions and the Expansion of Distribution Liability

The European Union represents the strongest example of regional media restriction in the RT case. After Russia's escalation of the war in Ukraine, European policy shifted from treating RT primarily as foreign propaganda to treating its distribution as part of a hostile wartime

information environment. The key concept was not merely that RT could publish false or slanted content. The broader claim was that Russian state media functioned as an instrument of state power, coordinated messaging, and destabilizing influence. This changed the legal meaning of distribution.

In a traditional media-liberty framework, the usual response to propaganda is counter-speech: expose it, criticize it, label it, debate it, or contextualize it. The European response went further by restricting distribution of certain Russian state-media outlets. This move reflected a wartime judgment that the danger was not only individual false statements but the sustained circulation of an adversarial state's information system within the European public sphere. Under that logic, the regulated object is not simply one article, one broadcast, or one claim. It is the institutional channel and its capacity to reach audiences.

This has important consequences for copied articles. If the restriction is understood narrowly as a ban on official RT television broadcasts, then private reposting of text may appear outside the central concern. But if the restriction is understood broadly as preventing the availability or distribution of RT content to European audiences, then copied text becomes more legally and politically sensitive. The platform may ask whether allowing private users to repost RT content defeats the purpose of restricting RT's official channels. A copied article without a URL may still carry the content the restriction was designed to limit.

The expansion of distribution liability is one of the defining features of the European approach. Distribution no longer means only a television signal, satellite feed, or official website. It can include internet platforms, apps, search access, social-media accounts, embedded videos, mirrors, and potentially other means by which restricted content remains available to the public. The more expansive the definition becomes, the more difficult it is for platforms to draw a clean line between original publication and secondary circulation.

This does not mean that every private discussion of RT is prohibited or that every quotation is equivalent to operating RT's channel. Democratic societies still require the ability to discuss propaganda, quote it, refute it, study it, archive it, and analyze it. A scholar, journalist, or citizen may need to reproduce limited portions of RT material to explain what was said and why it matters. The difficult question is where quotation ends and redistribution begins. A short excerpt embedded in criticism is one thing. A full article copied daily with minimal commentary is another. The first use points beyond the source; the second may function as a replacement distribution channel.

The European model also reveals how the legal status of a post can depend on audience location rather than author location. A U.S.-based page may be lawful from the author's point of view, but if it makes RT content available to European viewers, the platform may regard that

availability as a European compliance issue. Facebook, not the author, is the actor most exposed to regional enforcement. Therefore the platform has an incentive to over-restrict, localize, or demote content before legal conflict arises. This is how regional law can shape global speech without directly prosecuting the individual speaker.

The practical result is partial visibility. The author may not be banned. The post may remain available in the United States. The same post may be unavailable, less visible, or functionally undiscoverable inside parts of Europe. In this model, the platform does not need to announce a universal rule. It can silently harmonize the post with European risk. That is the essence of legal geography: the same digital object enters different legal states depending on where it is viewed.

4.2. The United Kingdom and Canada: Broadcast Regulation Versus General Web Access

The United Kingdom and Canada illustrate a different pattern. Their restrictions on RT have been strongly tied to broadcast regulation and authorized distribution rather than a simple nationwide prohibition on every private act of reading, quoting, or discussing RT material. This distinction is important because broadcast systems have long been governed differently from individual speech. A broadcaster may require a license. A cable or satellite distributor may need authorization. Regulators may impose standards concerning accuracy, harm, impartiality, or public interest. Removing a channel from authorized distribution is not identical to making all discussion of that channel illegal.

In the United Kingdom, the revocation of a broadcast license signals a judgment about the fitness of a broadcaster to operate within the regulated media system. It does not necessarily mean that a private person cannot mention, criticize, quote, or analyze that broadcaster's material. Similarly, in Canada, removing a service from the list of authorized distribution options affects cable, satellite, and broadcasting systems. It does not automatically create a universal web firewall around every sentence produced by the outlet. The regulatory center is distribution by licensed or authorized systems, not every possible private reference.

This creates a more ambiguous environment for Facebook posts. If a U.S. user copies an RT article, a viewer in the United Kingdom or Canada may not be encountering RT through a licensed broadcast channel. They are encountering user-generated content on a global platform. The post is not the same object as an RT television signal. But if the post reproduces RT content in full, it may still raise the question of whether the platform is enabling the outlet's restricted material to circulate through another route. The legal answer may differ from the policy answer. A platform may restrict content even where the law does not clearly require full removal.

Broadcast regulation also shows how the same outlet can be restricted in one medium but accessible in another. A channel may be removed from television packages while its articles remain reachable online. Its official account may be banned from one platform but not another.

Its clips may be unavailable on one service while copied text circulates in private groups. Therefore saying “RT is banned in the UK” or “RT is banned in Canada” is imprecise unless the medium is specified. Banned as a licensed broadcaster? Removed from cable distribution? Blocked as a website? Restricted as a Facebook link? Suppressed as copied text? These are different questions.

PressTV has a comparable but not identical history in broadcast regulation. The outlet lost licensed access in some contexts and has faced sanctions and distribution disruptions. Yet the same principle applies: removing a channel from regulated broadcasting does not necessarily erase every reference to its content. A copied PressTV article on Facebook may not be treated as a broadcast, but it may still be interpreted as republication, state-media amplification, or a sanctions-sensitive object depending on the platform’s policy and the viewer’s jurisdiction.

The United Kingdom and Canada therefore show a middle position between American openness and broader wartime regional restrictions. They may restrict official distribution channels without imposing a total information blackout. For the Facebook user, this means that the post might remain visible to UK or Canadian viewers even if the official broadcaster is not authorized in normal broadcast systems. But visibility is not guaranteed. Platform policy can be stricter than broadcast law. Local-law complaints can still occur. Copyright claims can still remove full-text reposts. User reports can trigger review. Algorithmic ranking can still contain the post.

This distinction helps avoid overstatement. Not every restriction is a total ban. Not every broadcast ban is a web block. Not every web block reaches copied text. Not every copied text post is safe from platform action. The modern environment is layered, and the UK-Canada pattern is one of partial regulatory exclusion rather than simple national erasure.

4.3. Australia, Sanctions, and Reputational Enforcement

Australia illustrates the role of sanctions and reputational enforcement in shaping visibility. Sanctions do not always operate like website blocks. They often target persons, entities, transactions, assets, services, financial relationships, and institutional dealings. A sanctioned media organization may become difficult to pay, host, advertise, distribute, partner with, or platform. But sanctions may not necessarily create a clear rule that every private citizen is forbidden from reading or discussing the entity’s public statements. The ambiguity lies in the space between information and support.

PressTV is particularly relevant here because restrictions around Iranian state-linked media often arise through sanctions and human-rights concerns rather than a single media-ban structure. If a state designates a media outlet or associated institution, companies may become cautious. They may remove accounts, block monetization, avoid carrying feeds, end advertising relationships, or prevent transactions that could be interpreted as providing services. A platform

may choose to restrict content not because every post is illegal, but because the source has become a compliance and reputational risk.

Reputational enforcement can be as powerful as legal enforcement. A platform does not want to be accused of amplifying sanctioned state propaganda, hosting forced-confession media, spreading foreign influence operations, or serving as a workaround for restricted outlets. Even if a particular user's post is legally defensible, the platform may prefer a conservative policy. It may remove official accounts, limit discoverability, reduce recommendation, restrict monetization, or apply labels. These measures may be justified publicly as authenticity protection, safety, civic integrity, sanctions compliance, or misinformation reduction.

Australia also shows how national policy can influence platform behavior beyond the formal boundaries of law. A company operating globally may create internal risk categories for state-linked outlets. Once a source is associated with sanctions, human-rights abuses, disinformation campaigns, or foreign interference, the platform may treat the source as high-risk everywhere. The result is not necessarily a visible Australian block. It may be a global platform posture shaped partly by Australian, European, American, Canadian, and British pressure together.

This is how reputational enforcement produces invisible boundaries. The platform may not tell the user, "Your post was reduced because it reproduces content from a sanctioned Iranian state-media outlet." It may simply show the post to fewer people. It may decline to recommend it. It may prevent advertising. It may place the account in a lower-trust category. It may wait for repeat behavior before acting. Or it may do nothing because the post is too small to matter. The user cannot easily distinguish these outcomes from ordinary low engagement.

Sanctions also interact with copyright and platform terms. A full copied article can be challenged as unauthorized republication even when the political content itself is not removed. This creates a convenient nonpolitical enforcement pathway. A platform may not need to decide whether a user's PressTV repost is propaganda. It can remove the post if it violates copyright or terms against repeated copying. Conversely, if the outlet itself is unlikely to file copyright complaints, the copied article may remain. Enforcement may therefore depend less on principle than on who complains, what systems detect, and whether the post attracts attention.

Australia's role in this paper is to demonstrate that legal geography is not only about hard blocking. It is also about risk atmospheres. A country can contribute to the global restriction environment by sanctioning an entity, condemning an outlet, pressuring platforms, joining allied policy statements, or defining a source as part of a hostile influence ecosystem. Those signals travel through corporate compliance departments and platform policy teams. The user never sees that process. He sees only the final interface.

For a U.S.-based Facebook user, the Australian example means that foreign visibility cannot be inferred from domestic publication. Even if Australia does not create a direct technical block of a particular post, its sanctions and policy stance may contribute to broader platform caution around PressTV or similar outlets. The post may exist, but the risk category around it has changed.

4.4. Local Law Restriction as a Country-Specific Visibility Layer

Local-law restriction is the mechanism that most clearly explains how one post can remain visible in one place and disappear in another. A platform may determine that content violates the law of a particular country but not its global community standards. Instead of deleting the post everywhere, the platform restricts access only in the jurisdiction where the legal issue arises. To the author, especially if he is outside that jurisdiction, the post may appear normal. To the affected viewer, it may be missing, replaced by an unavailable-content notice, or absent from feeds and search.

This mechanism is designed to solve a difficult corporate problem. A global platform cannot easily maintain one universal speech rule acceptable to every country. Germany, India, Turkey, Brazil, the United States, Canada, Australia, the European Union, and the United Kingdom do not regulate speech in the same way. Defamation, hate speech, extremist content, election law, privacy, sanctions, state media, and national security all vary. If a platform removed every locally unlawful post globally, the most restrictive jurisdiction could determine worldwide speech. If it ignored local law, it could face fines, blocking, criminal exposure, office raids, app restrictions, or political retaliation. Country-specific restriction is the compromise.

For the user, this compromise produces opacity. The author may not know that a post is restricted elsewhere. The dashboard may not show country-level reach. The post may continue to receive domestic engagement. Friends in the United States may still comment. The author may conclude that no restriction exists. But a foreign viewer may see nothing. If the author lacks contacts in the affected country or lacks analytic visibility by country, the restriction can remain undetected indefinitely.

Local-law restriction is particularly relevant for RT because some jurisdictions treat RT distribution as legally restricted. If Facebook determines that a user's copied article functions as RT content available to viewers in those jurisdictions, the platform could theoretically limit visibility there while leaving the post available elsewhere. Whether it does so consistently is a separate question. Large platforms are not perfectly uniform. They may catch some posts and miss others. They may restrict official accounts but not copied text. They may rely on reports. They may change policy after court rulings, political pressure, or enforcement actions. Inconsistency is part of the system.

The same mechanism can apply to other categories of content. A post may be lawful in the United States but violate privacy law in another country. It may be protected political opinion in one place and unlawful extremist support in another. It may be ordinary criticism in one jurisdiction and prohibited defamation in another. It may be acceptable commentary under one copyright regime and infringing under another. The platform's country-specific visibility layer allows the same post to have different legal states in different places.

This layer also interacts with the user's proposed VPN test. A VPN can sometimes reveal crude geoblocking by making the user appear to come from another country. But it cannot fully reproduce the experience of a normal foreign viewer. Facebook may treat logged-out users differently. It may use account region, device history, language, identity, social graph, and behavioral signals. A newly created foreign account may be flagged as suspicious. A VPN may reveal one layer of location-based availability while missing algorithmic distribution, feed ranking, or account-specific treatment. Human verification by real viewers in the target country is therefore more meaningful.

Country-specific restriction also changes the meaning of censorship. Under the old model, censorship was a universal denial: the page would not load. Under the new model, censorship may be a local absence inside a global interface. The author sees the post. Some readers see the post. Other readers never encounter it. The post exists, but not everywhere. It is public, but not for every public. It is published, but legally partitioned.

This is the deepest implication of legal geography. The digital public sphere is no longer one space. It is a mosaic of overlapping national permissions, corporate risk calculations, platform rules, and algorithmic choices. The same Facebook post can be speech, evidence, propaganda, copyright infringement, sanctioned-media distribution, protected commentary, or legally restricted content depending on where it lands and how it is framed. The user does not publish into the world. He publishes into a machine that rebuilds the world differently for each viewer.

5. Platform Governance: Facebook as a Private Border System

Facebook is often described as a social network, but for political information it functions more like a private border system. It decides what may enter, what may circulate, what must be labeled, what must be removed, what may be monetized, what may be recommended, and what may cross from one audience to another. It does not merely host posts. It sorts them. It weighs them. It classifies their sources. It assigns risk. It interprets behavior. It localizes access by country. It decides whether content remains public, becomes obscure, or disappears.

This makes Facebook different from the open web. On the open web, a user may reach a source directly through a browser, search engine, bookmark, mirror, or archive. On Facebook, the post

enters a managed environment. The platform controls the interface, the feed, the recommendation system, the share mechanism, the advertising system, the reporting system, the copyright system, the page dashboard, and the country-specific visibility layer. A post that appears successfully on a page has not necessarily passed through all these gates. It has only entered the first room.

The platform is therefore a borderless territory with many borders. Some borders are legal. Some are technical. Some are commercial. Some are reputational. Some are algorithmic. Some are invisible. A national government may block a website openly, but a platform can reduce distribution without announcing the reduction. A court may order removal in a country, but a platform can also act preemptively to avoid legal risk. A user may violate no criminal law while still triggering a platform's state-media policy, spam policy, copyright policy, authenticity policy, or recommendation-quality system. The outcome may be removal, but more often it may be containment.

For RT and PressTV, this platform-border function is central. These outlets are not merely domains. They are institutional signals. Their names carry legal, geopolitical, and reputational meaning. A Facebook post containing their copied text may be interpreted through the platform's state-media framework even when no URL appears. A repeated pattern of posting full articles may be interpreted differently from occasional quotation. A page that reposts foreign state-media material daily may be evaluated not only post by post but as an account-level behavior pattern. The platform may ask not only, "What does this sentence say?" but "What is this page doing?"

This transforms the meaning of publication. The user may believe he has placed an article before the public. Facebook may instead treat the article as an object that should be shown narrowly, kept away from recommendations, restricted in certain jurisdictions, excluded from advertising, or allowed only because it has not attracted enough attention. The post remains, but the border system has narrowed its travel.

5.1. Meta's State-Media Policies and Entity-Based Enforcement

State-media policy differs from ordinary content moderation because it focuses on the identity and control structure of the source. Ordinary moderation usually asks whether a post contains prohibited material: threats, harassment, sexual exploitation, fraud, spam, incitement, graphic violence, or other disallowed content. State-media enforcement asks a different question: who stands behind the outlet, and what role does the outlet play in the information strategy of a state?

This shift from content to entity is crucial. A state-media article may contain accurate facts, ordinary reporting, or legitimate criticism of Western policy. Yet the platform may still classify

the source as state-controlled, state-funded, state-directed, or state-aligned. The classification does not necessarily declare that every article is false. It declares that the source operates within a political and institutional structure that affects how its content should be distributed, labeled, or trusted.

Entity-based enforcement is especially important for RT. A platform does not have to prove that every RT article is false before restricting RT as an outlet. It can decide that RT as an entity belongs to a prohibited or high-risk class. Once that decision is made, enforcement can attach to the outlet's accounts, pages, channels, domains, apps, videos, advertising privileges, and potentially recognizable reposted material. The platform's concern becomes the survival of the outlet's distribution network, not the truth status of each individual sentence.

PressTV presents a related case. It may be treated as Iranian state media, a sanctions-sensitive entity, a source associated with state narratives, or an outlet whose official platform presence creates compliance and reputational risks. The precise legal and policy treatment may differ from RT, but the logic of entity recognition remains. The platform does not only see an article. It sees a source identity.

For the private user, entity-based enforcement creates ambiguity. The user is not RT. The user is not PressTV. The user may be an American citizen posting in his own name, with his own commentary, for his own reasons. But the material being copied may still originate from an entity the platform restricts or distrusts. The post becomes a hybrid object: private speech carrying state-media content. That hybrid status creates the possibility of inconsistent treatment. The user may be left alone. The post may be demoted. The page may be watched. The content may be restricted only in some countries. The platform may act only if the pattern becomes large, systematic, or reported.

Entity-based enforcement also makes the removal of URLs less decisive. If the platform's concern were only domain linking, removing the link would solve much of the problem. But if the concern is the circulation of an entity's content, then copied text can still matter. A platform may identify the article by text matching, title matching, language patterns, repeated posting behavior, user reports, or comparison with known restricted sources. The link is only one signal among many.

This is why the absence of trouble over a year should be interpreted carefully. It may mean the platform has not classified the behavior as serious. It may mean the page is small enough not to matter. It may mean copied text is harder to detect than links. It may mean the posts are mostly treated as user commentary. It may mean the posts are visible domestically but contained internationally. It may mean no one has reported them. It does not prove that the material is fully distributed or immune from future enforcement.

Entity-based enforcement changes the politics of speech by making provenance part of moderation. The words remain important, but the institutional identity behind them becomes a regulatory object. A copied RT article is not just an article. A copied PressTV article is not just an article. On a platform organized around risk classification, source identity travels with the text even when the hyperlink is gone.

5.2. Link Blocking Versus Text Recognition

Link blocking is the simplest form of platform control. A post contains a domain. The platform recognizes the domain. The platform labels it, blocks it, reduces its reach, warns before opening it, prevents sharing, or removes the post. This kind of enforcement is fast, scalable, and easy to explain. It treats the URL as the object of control.

But link blocking is only the surface layer. Modern platforms can also analyze text, images, video, metadata, behavior, and networks. A copied article without a URL may still be identifiable through exact phrases, headlines, paragraph structure, distinctive names, source attributions, repeated posting patterns, or user reports. If the text was already indexed, archived, reported, or associated with a known source, the platform may be able to identify it even after the link has been removed.

Text recognition can serve many purposes. It can detect spam, copyright infringement, repeated misinformation, terrorist propaganda, coordinated campaigns, scam templates, and previously removed content. It can also detect near-duplicate posts that appear across multiple pages or groups. In the state-media context, text recognition can help identify content that has migrated away from official channels into secondary circulation. A platform concerned about RT or PressTV does not need to rely entirely on the presence of an RT or PressTV link.

At the same time, text recognition is imperfect. It may miss copied articles. It may fail when the article is edited, summarized, paraphrased, translated, truncated, or screenshot. It may not be deployed aggressively for all sources. It may detect but not enforce. It may require a report before review. It may be reserved for high-risk events such as wars, elections, pandemics, or sanctions escalations. It may operate differently by language and region. Therefore the survival of copied articles on a page does not prove that the platform cannot detect them. It only proves that detection has not produced obvious removal.

This imperfection creates the ambiguous zone in which many users operate. Direct links to restricted sources may be more likely to trigger labels or blocks. Copied text may pass through more easily. Screenshots may pass through more easily still. Summaries with commentary may be treated as original speech. Each transformation moves the post farther from the easily classified source object. But each transformation also changes the meaning of the user's act. Removing the link can make the source less transparent. Copying the whole article can look like

mirroring. Posting a short excerpt with commentary can look like criticism. The technical form affects the political interpretation.

For RT and PressTV, link blocking and text recognition create different risk profiles. A direct RT link may trigger immediate source-based treatment. A copied RT article may avoid domain-level enforcement but still carry recognizable content. A PressTV link may be treated according to sanctions, platform policy, or state-media labeling. A copied PressTV article may be less visible to automated systems but still subject to reports, copyright claims, or account-pattern review. The absence of a URL reduces one signal but may strengthen another: the appearance that the user is deliberately bypassing source restrictions.

The distinction also matters for readers. A link preserves the path to the source. It allows readers to see where the article came from, inspect the original page, compare context, and judge the outlet. Plain copied text can detach the article from its institutional setting. That detachment may be harmless when the source is clearly named and the user adds commentary. But it can be problematic when the copied text circulates as if it were ordinary user speech or neutral reporting. A platform concerned with foreign influence may view source detachment as part of the risk.

The copied-article problem therefore cannot be understood as a simple trick for avoiding link restrictions. It is a move from one enforcement layer to another. It may evade domain recognition, but it enters text recognition, copyright analysis, behavior analysis, and source-laundering suspicion. A post without a URL may be easier to publish and harder to classify, but it is not necessarily unrestricted.

5.3. Removal, Demotion, Warning Labels, and Reduced Distribution

Platform enforcement is often imagined as removal, but removal is only one outcome. A post can be left online while its circulation is reduced. It can receive a warning label. It can be excluded from recommendations. It can be made ineligible for advertising. It can be hidden from certain countries. It can be shown mostly to followers. It can be suppressed in search. It can be prevented from appearing in groups. It can be deprioritized in feeds. It can be limited without the author receiving a visible penalty.

This hierarchy of enforcement is central to soft censorship. Removal is visible and contestable. A removed post produces evidence. The author can complain, appeal, screenshot the notice, or publicize the action. Demotion is harder to detect. The post remains, but fewer people see it. The author may blame audience disinterest, timing, the algorithm, or random variation. Reduced distribution is therefore less politically costly for the platform than deletion. It controls reach while preserving the appearance of publication.

Warning labels occupy a middle position. They do not always remove the content, but they alter how it is received. A label may identify a source as state-controlled media, disputed, sensitive, or context requiring. Labels can reduce trust, reduce sharing, trigger user hesitation, or mark the post for further algorithmic handling. They also allow platforms to claim they are informing users rather than censoring them. For state-media content, labeling can be framed as transparency. For the speaker, labeling may function as a reputational penalty.

Demotion is more opaque. A demoted post may still be visible on the page but not widely delivered in feeds. It may be visible to people who visit the page directly but not recommended to others. It may reach existing followers but fail to reach non-followers. It may be excluded from certain ranking systems that normally produce expansion. In practice, demotion can be more consequential than removal for small accounts. A removed post is gone, but a demoted post creates the illusion that the audience simply did not respond.

Reduced distribution can also be source-specific, topic-specific, behavior-specific, or account-specific. A single post might be downranked because it includes copied state-media text. A page might be downranked because it repeatedly posts such material. A topic might be downranked during a crisis. A country-specific rule might reduce reach in one jurisdiction. A civic-content policy might reduce political distribution generally. A copyright or spam classifier might lower the post's ranking because it is duplicate text. These systems can overlap, making the cause impossible for the user to identify.

RT and PressTV content may encounter several of these mechanisms simultaneously. A direct RT link may be blocked or labeled. A copied RT article may be demoted if recognized. A PressTV repost may be left visible but not recommended. A full article may be treated as low-originality duplicate content. A daily pattern may be treated as spam-like. A foreign viewer may be blocked by local-law restrictions. The user sees only the final surface: the post is up.

The dashboard may not reveal the difference between ordinary low reach and enforced low reach. It may show impressions, likes, comments, or reach without explaining whether distribution was reduced. It may not show country-level visibility. It may not show how many people would have seen the post absent demotion. It may not distinguish follower reach from recommendation reach in a way sufficient to diagnose containment. The user is left to infer from patterns.

This is why comparison is necessary. One post teaches little. A pattern can teach more. If ordinary posts consistently reach non-followers and foreign viewers while copied RT or PressTV posts remain inside the existing audience, that suggests containment. If RT posts underperform PressTV posts, that suggests source-specific treatment. If full copied articles underperform short excerpts with commentary, that suggests form-specific treatment. If posts are visible to U.S.

viewers but unavailable to European viewers, that suggests legal geography. Without such comparisons, the author's experience remains anecdotal.

The most important point is that platform enforcement is no longer binary. The post may not be alive or dead. It may be alive but weak. It may be public but not distributed. It may be legal but not recommended. It may be visible but not discoverable. It may be present but socially absent. Removal is only the most obvious border. Reduced distribution is the more modern one.

5.4. Why Enforcement Can Be Real but Inconsistent

Large platforms enforce rules unevenly because their systems are enormous, automated, multilingual, politically contested, and constantly changing. Inconsistency is not an exception to platform governance. It is one of its normal features. A post can violate a rule and remain up. Another post can be removed mistakenly. A direct link can be blocked while copied text survives. One country can see a restriction before another. One page can be ignored because it is small while another is penalized because it is influential. One reviewer may interpret a post as commentary; another may interpret it as amplification.

This inconsistency leads users to draw the wrong conclusions. If a post remains up, the user may conclude that the platform permits the content. If similar posts remain up for a year, the user may conclude that there is no restriction. But platforms often enforce at thresholds. They may act only when content is reported, when an account grows, when a topic becomes politically sensitive, when a legal demand arrives, when a court ruling changes risk, when an automated classifier is updated, or when a source is newly designated. Past survival is evidence of past non-removal, not proof of future safety or broad visibility.

Inconsistent enforcement can also be strategic. Platforms may avoid high-profile removals when removal would generate accusations of censorship. They may prefer demotion because it is quieter. They may treat small accounts leniently because they pose little distribution risk. They may restrict official state-media accounts but not private discussion because the latter is harder to distinguish from legitimate commentary. They may enforce heavily during elections or wars and lightly during quieter periods. They may act differently after government pressure. The same written rule can produce different outcomes across time.

Human review adds another layer. A reviewer may see a copied RT article and ask whether the user is criticizing it, archiving it, endorsing it, or mirroring it. If there is commentary, the post may be allowed. If there is no commentary, it may look like redistribution. If the page posts such articles daily, the pattern may matter more than the individual post. If the article contains additional policy-violating claims, it may be removed for reasons unrelated to RT. If the article is merely geopolitical reporting, it may remain. Review depends on context, and context is expensive to evaluate at scale.

Automation adds still more variation. Automated systems are better at detecting links, known media hashes, repeated text, and high-confidence violations than at interpreting intent. A copied article may be detected as duplicate content but not as state media. It may be detected as state media but allowed because the account is not official. It may be detected only if many users post the same article. It may be missed because of formatting changes. It may be classified differently after translation. Thus the same behavior may produce no visible action for long periods.

Legal changes can also alter enforcement suddenly. A court ruling, sanctions designation, regulatory demand, or government request can shift the platform's risk calculation. A pattern tolerated yesterday may become more sensitive tomorrow. This is especially true for wartime information policy. RT, PressTV, and similar outlets exist within geopolitical conflicts. As those conflicts change, platforms may tighten or loosen controls. The user's year of trouble-free posting may reflect a period of low enforcement rather than a stable right of distribution.

Inconsistency should therefore be understood as part of soft censorship's design environment. The user cannot easily know whether he is free, ignored, tolerated, demoted, or contained. The platform may not tell him. The audience may not know what it is missing. Regulators may see only official compliance reports. Researchers may need external testing to detect regional restrictions. The system produces ambiguity as a byproduct of scale and as a convenience for power.

For the U.S.-based Facebook user posting RT and PressTV articles, the most reasonable conclusion is cautious. The absence of overt punishment is real evidence that the account is not under obvious removal enforcement. It is not evidence that the posts are fully distributed. It is not evidence that foreign viewers can see them. It is not evidence that Meta treats full copied articles the same as original commentary. It is not evidence that enforcement will remain unchanged. The platform border system may already be operating, but operating quietly.

Facebook as a private border system therefore changes the nature of the public square. A citizen can still stand and speak, but the platform controls the doors, corridors, loudspeakers, maps, and exits. The speaker may not know which doors are locked. He may only know that he can still hear himself.

6. The Copied-Article Problem

The copied-article problem sits at the intersection of speech, attribution, copyright, sanctions, platform policy, and political meaning. A user who copies an article from RT or PressTV into a Facebook post without including the original URL may think he has changed the nature of the object. The post no longer contains a direct link to a restricted or disfavored source. It appears

as ordinary user text. It may be surrounded by the user's own framing, placed into a personal or public page, and received by followers as part of a continuing political conversation. From the author's point of view, the act may feel like sharing information, preserving a record, comparing narratives, or bypassing what he regards as unfair exclusion from mainstream discourse.

But the platform, the law, and the audience may see a different object. The copied article remains a text produced by a particular institution. Its institutional origin may matter. If copied in full, it may function as republication rather than commentary. If copied every day, it may function as a substitute distribution system. If copied without enough source identification, it may detach state-media content from the state-media label that platforms or readers would otherwise see. If copied without transformation, it may create copyright exposure. If copied from a sanctioned or restricted outlet, it may become part of a compliance problem even if the user is not formally connected to the outlet.

The copied-article problem is therefore not solved by removing the URL. The URL is only one sign of origin. Source identity can remain visible in the headline, style, phrasing, named outlet, article structure, embedded claims, repeated subject matter, and posting pattern. Even when a platform does not detect the source automatically, human readers may. Even when readers do not, the act of full-text copying changes the legal and ethical position of the post. The user is no longer merely pointing to a source; he is reproducing the source.

This distinction matters because modern platforms often treat linking, excerpting, quoting, summarizing, screenshotting, and full reposting differently. A link says, "Go there." A quotation says, "Consider this passage." A summary says, "Here is the argument in my words." A screenshot says, "Here is a record of what appeared." A full copied article says, "Here is the article itself." Each form has a different relationship to the original publisher, to the audience, and to the platform's risk systems. The same political content may be tolerated in one form and restricted in another.

For RT and PressTV, the copied-article problem is especially sharp because the content is politically charged and institutionally marked. The issue is not only whether the words are true or false. The issue is whether the user's page becomes a relay point for outlets that may be restricted, demoted, sanctioned, or reputationally disfavored elsewhere. A platform may not care if a user criticizes RT. It may care if the same user reproduces RT articles daily in a way that preserves RT's distribution while stripping away some of the platform's normal source controls. Similarly, a PressTV article may be treated differently when discussed as evidence of Iranian state framing than when copied wholesale as ordinary news.

The copied-article problem therefore forces a practical distinction between publication and mediation. The user is not simply sharing content. He is mediating between a restricted or

controversial source and an audience. The form of that mediation determines whether the act looks like criticism, research, archiving, republication, or circumvention.

6.1. Full Reposting Versus Quotation and Commentary

Full reposting is the most exposed form because it reproduces the original article as a substitute for the original publication. The reader does not need to visit RT or PressTV. The article has been transplanted into Facebook. If the post contains the full headline, full body, and minimal original analysis, the user's page becomes a mirror of the source. It may not be an official mirror, and it may not be coordinated with the outlet, but functionally it performs a similar act: it preserves distribution when direct distribution may be restricted, demoted, or reputationally discouraged.

Quotation and commentary are different. A user who quotes a limited passage and then analyzes it is not simply replacing the original outlet. He is using the article as an object of discussion. The quoted material becomes evidence for an argument. The user's own contribution is the reason for the post. The source is present, but it is subordinate to criticism, comparison, explanation, or interpretation. This is the traditional mode of journalistic, scholarly, and civic engagement with objectionable or controversial material. One must be able to quote propaganda in order to analyze propaganda. One must be able to cite an adversary's claim in order to refute it.

The difference lies partly in quantity but more deeply in function. A long excerpt may be defensible if it is necessary to analyze a specific claim, rhetorical structure, or evidentiary pattern. A short excerpt may still be problematic if it is repeatedly used as a teaser to drive traffic or reproduce a source's message without scrutiny. Full reposting is riskier because it generally supplies the reader with the whole substitute object. Commentary is safer because it transforms the object into something else: analysis, criticism, documentation, or argument.

Frequency also matters. One copied article may be a data point. Daily copied articles become a pattern. A pattern can transform interpretation. The user may see each post as independent, but the platform may see a recurring behavior: a page repeatedly republishing state-linked media content. Even if each article is permitted individually, the aggregate pattern may be interpreted as systematic amplification. This is especially true when the user posts the article without substantial commentary, without comparison to other sources, and without clear framing that explains why the article is being reproduced.

Attribution matters as well. Naming the source is ethically important because readers should know where the material came from. But attribution does not by itself solve the problem. A full copied article with attribution can still be unauthorized republication. It can still function as a mirror. It can still carry restricted source content into jurisdictions where the platform would otherwise limit the outlet. Conversely, failing to name the source creates another problem: the

reader may not understand that the material originated from a state-linked outlet. That can make the post appear less transparent and more like source laundering.

The strongest form, analytically and ethically, is therefore not silent full reposting but framed quotation. The user identifies the source, quotes only what is needed, explains why it matters, compares it with other reporting where appropriate, and adds independent judgment. This does not guarantee platform distribution, but it changes the character of the act. The post becomes the user's speech about RT or PressTV rather than RT or PressTV's speech reproduced by the user.

This distinction is central to soft censorship because platforms are more likely to tolerate discussion than substitution. A platform can defend allowing citizens to analyze hostile state media. It has a harder time defending a system that lets users recreate the distribution networks of restricted outlets through full-text reposting. The boundary is not always clear, but the direction is clear: the more original commentary dominates the post, the less the post resembles a mirror.

6.2. Copyright Risk Independent of Political Content

Copyright is separate from political restriction. A full article can create copyright risk regardless of whether the source is RT, PressTV, the New York Times, the BBC, Al Jazeera, Fox News, Reuters, or a local newspaper. The political status of the source does not erase the fact that articles are protected expressive works. Copying an entire article into Facebook may infringe copyright even if the user gives credit, includes no URL, earns no money, and intends to inform rather than profit.

This matters because copyright provides a nonpolitical route for removal. A platform does not need to decide whether a copied RT article is propaganda. It can remove the post if the copyright holder complains. It does not need to decide whether a PressTV article is sanctioned state media. It can remove the post if it is unauthorized republication. Copyright enforcement can therefore produce a result that feels like political censorship even when the formal reason is intellectual property. Conversely, if the source does not complain, full reposts may remain online for long periods. The absence of removal may simply mean the copyright owner has not enforced its rights.

Fair use or analogous quotation principles may protect some uses, especially criticism, commentary, teaching, scholarship, research, news analysis, or parody. But fair use is not a magic label. It depends on purpose, transformation, amount used, and market effect. Copying an entire article is harder to defend than quoting limited passages. Adding commentary helps because it transforms the use from substitution toward analysis. Using the article to criticize the source or compare narratives is stronger than merely reposting it as news. The more the user's

original expression dominates, the stronger the argument that the post is commentary rather than replacement.

The market effect factor is particularly relevant for full reposting. If the Facebook post gives readers the whole article, it may substitute for visiting the original source. In ordinary media copyright analysis, that weighs against the copier. With RT or PressTV, the situation is unusual because the source may not object to wider circulation. A state-media outlet may welcome republication. But copyright law does not disappear merely because the publisher benefits politically. Also, platforms cannot build moderation systems on the assumption that every publisher welcomes copying. They need general rules, and general rules often disfavor full-text reproduction.

Copyright also interacts with platform ranking. Even absent a formal copyright complaint, platforms may downrank duplicate or low-originality content. A full copied article may be treated as less original than a post containing the user's own analysis. It may resemble spam, scraping, or automated content. This means full reposting can reduce distribution even when no legal complaint occurs. The post may survive but perform poorly because the platform prefers original user expression over copied publisher text.

There is also an evidentiary problem. If a user copies entire articles daily, the page archive itself documents the pattern. If a copyright complaint or platform review later occurs, the user cannot easily claim isolated quotation. The repeated structure of the posts may show that the page functions as a republication channel. Even if the user's motive is political comparison, the form may tell another story.

This is why copyright should be treated as independent risk. A user may believe the central issue is censorship of RT or PressTV, but the first actual enforcement action may come through copyright, duplicate-content policy, or spam-like behavior. Political restriction and copyright restriction can converge on the same post for different reasons. The user experiences one removal; the system may have many possible justifications.

The practical lesson is that quotation plus commentary is not only stronger politically; it is stronger legally. It preserves the ability to discuss restricted state media while reducing the appearance of unauthorized republication. It places the user's own speech at the center. It makes the source an object of analysis rather than a substituted publication.

6.3. The Difference Between Archiving, Criticism, and Acting as a Mirror Site

A copied article can serve several purposes, and those purposes matter. Archiving preserves material for future reference. Criticism interrogates the material. Acting as a mirror site redistributes the material. These categories can overlap, but they are not the same.

Archiving is concerned with preservation. A user may copy an article because he fears it will disappear, be altered, be blocked, or become difficult to find. This motive is common in contested information environments. People archive government statements, propaganda broadcasts, deleted posts, war claims, legal notices, and media narratives because future analysis requires a record. Archiving is especially important when platforms remove content or when governments restrict access. Without archives, the public record can become unstable.

But archiving has a form problem. A private archive kept for research is different from a public Facebook page that republishes entire articles daily. The first primarily preserves; the second distributes. Once an archive is public, searchable, shareable, and integrated into a social platform, it becomes more than storage. It becomes publication. If the archive is selective and repeatedly copies one state-media source, it may look less like neutral preservation and more like a relay system. The user's intent may be archival, but the public function may be distribution.

Criticism is different. It uses the source as an object to be examined. A critical post may quote a headline, reproduce a claim, compare it with other accounts, examine omissions, identify framing, or explain why the outlet's narrative matters. Criticism can be hostile, sympathetic, comparative, or skeptical. What defines it is transformation. The post does not simply carry the article forward; it changes the reader's relationship to the article. It asks the reader to interpret.

Acting as a mirror site is different again. A mirror preserves access to the original content in substitute form. It is often used when the original site is blocked, overloaded, censored, or taken down. In technical culture, mirrors can be legitimate tools of resilience. In platform governance, however, mirroring restricted content can be treated as circumvention. If RT's official distribution is restricted in a jurisdiction, a Facebook page that republishes RT articles in full may be interpreted as recreating part of that distribution. If PressTV's official accounts or domains are restricted, repeated full-text reposting may be seen similarly.

The difference often turns on the user's original contribution. If the page contains mostly copied articles, the page looks like a mirror. If it contains excerpts embedded in commentary, the page looks like analysis. If it contains copied articles from many sources for comparative study, it looks more archival or research-oriented. If it contains only one restricted source repeatedly, it looks more like amplification. If the user adds dates, source labels, commentary, and comparison, the archival function becomes more credible. If the user posts plain full text without context, the mirror function dominates.

The distinction also affects audience interpretation. Readers encountering full copied articles may receive the source's narrative directly. Readers encountering commentary are invited to evaluate the narrative. A platform concerned with foreign influence cares about this difference

because influence operations often depend on repeated exposure, not merely persuasion by one article. Mirroring increases exposure. Criticism can reduce influence by contextualizing the source. Archiving can go either way depending on how it is structured.

This does not mean a user must never reproduce a full article. There may be cases where full reproduction is necessary for documentation, especially if the article is likely to vanish and has public importance. But such cases are stronger when the user clearly states the archival purpose, adds analysis, limits frequency, avoids creating a substitute news feed, and preserves source identity. The more a page resembles a daily replacement distribution channel, the weaker the archival argument becomes.

For RT and PressTV, the safest conceptual posture is neither silent mirroring nor hidden copying. It is transparent critical mediation. The user can say, in effect: this is what this outlet is claiming; here is the relevant passage; here is why it matters; here is what other sources say; here is my judgment. That posture preserves the right to discuss disfavored sources while avoiding the strongest indicators of circumvention.

6.4. How Platform Systems May Interpret Repeated Verbatim Reposting

Platforms do not interpret posts only as isolated speech acts. They interpret patterns. Repeated verbatim reposting can trigger several possible classifications at once: duplicate content, low-originality content, spam-like behavior, copyright risk, state-media amplification, coordinated narrative distribution, or attempted circumvention of link restrictions. The platform may not need to know the user's motive. It can classify the behavior by its observable structure.

The key signals include frequency, completeness, source concentration, repetition, absence of commentary, engagement pattern, and audience expansion. A page that occasionally quotes RT or PressTV as part of broader political analysis looks different from a page that posts full articles from those outlets every day. A page that compares Russian, Iranian, American, British, and Arab sources looks different from a page that republishes one outlet. A page that adds substantial original analysis looks different from a page that pastes article text with little framing. A page whose posts are shared by a diverse audience looks different from a page circulating inside a narrow ideological cluster.

Automation may detect repeated copying even if it does not identify every source. Duplicate text itself can be a signal. Platforms often prefer original content because copied material can degrade feed quality, crowd out users' own expression, or indicate scraping and spam. A full repost may therefore be downranked simply because it is not original. If the same page repeatedly posts copied articles, account-level distribution may be affected. The user may not receive a warning because low-originality demotion is not always treated as a punishable violation. It may be treated as a ranking decision.

State-media amplification adds another layer. If the copied text is recognized as RT or PressTV, the platform may interpret repetition as a source-circulation pattern. The account does not need to be official. It may still function as a relay. The platform's concern may be especially strong if official accounts for the source are restricted, because secondary reposting can preserve distribution after official channels are reduced. This is where "no URL" can become suspicious. Removing the link may not look like neutrality; it may look like an attempt to move the same content through a less detectable form.

Human reports can activate review. A post may remain untouched for months until someone reports it. A critic, foreign viewer, political opponent, or copyright holder may flag the content. Once reviewed, the page's pattern may matter. The reviewer may see not one post but many. A page history showing daily full-text reposts from restricted state media may be more vulnerable than a single isolated article. The enforcement decision may be framed as spam, intellectual property, state-media policy, local law, or repeated violation of platform norms.

Country-specific systems may interpret the same pattern differently. In the United States, the posts may remain visible. In the European Union, RT-derived content may carry greater compliance risk. In the United Kingdom or Canada, the platform may distinguish broadcast restrictions from user discussion but still apply private policy. In Australia, sanctions and reputational concerns may shape platform caution around Iranian state media. A repeated pattern increases the chance that the post will encounter at least one restrictive layer somewhere.

There is also the issue of account reputation. Platforms maintain internal trust signals. Although users rarely see them directly, account behavior can affect distribution. Repeated posting of copied political content, especially from sensitive state-linked sources, may lower trust or reduce eligibility for recommendation. The page may not be banned; it may simply stop growing. Existing followers may still see posts, while non-followers do not. The author may interpret this as a stable but small audience. In fact, it may be containment.

The user's observation that he has had no problem for a year should be entered into the analysis as evidence but not as proof. It shows that repeated verbatim reposting has not led to obvious deletion, suspension, or warning. It does not show that the posts are treated normally by ranking systems. It does not show that the posts reach foreign viewers. It does not show that non-followers see them. It does not show that future policy changes will leave them untouched. It does not show that a human review would interpret the pattern favorably.

The practical implication is straightforward. If the goal is to preserve broad visibility and reduce enforcement risk, repeated verbatim reposting is the weakest form. Excerpts plus commentary are stronger. Comparative analysis is stronger still. A post that clearly belongs to the user

intellectually is less likely to be treated as a mirror. A post that merely moves another outlet's article into Facebook is more likely to be treated as copied content, state-media relay, or restricted-source circumvention.

The copied-article problem therefore reveals the core logic of modern platform governance. The question is not only what the user says. It is what the user's page does. Does it analyze? Does it archive? Does it criticize? Does it compare? Or does it reproduce? Platforms may not answer that question fairly or consistently, but they are likely to ask it. And if they ask it, the form of the post becomes as important as its content.

7. Siloing as a Form of Soft Censorship

Siloing is the condition in which speech remains technically present but socially, geographically, or algorithmically enclosed. The speaker is not banned. The post is not necessarily removed. The account may remain in good standing. The archive may continue to grow. Yet the speech does not move freely through the public sphere. It circulates inside a bounded space: a known audience, a permitted jurisdiction, a narrow recommendation channel, or a politically compatible network. The speaker continues to speak, but the audience is quietly managed.

This is why siloing must be understood as a form of soft censorship. It does not resemble book burning, broadcast jamming, police seizure, or a browser error message. It resembles normalcy. The author sees the post. Friends react. A few comments appear. The dashboard shows activity. Nothing announces that a boundary has been reached. But the post may not reach non-followers, may not appear in search, may not cross into certain countries, may not be recommended, may not be shareable by some viewers, or may not travel beyond an already sympathetic circle. The public square remains visible from inside the room, but the doors have narrowed.

Soft censorship works by altering probability rather than possibility. Hard censorship says, "This cannot be seen." Soft censorship says, "This may be seen, but rarely, locally, and by the already initiated." It reduces reach instead of eliminating existence. It changes distribution instead of publication. It converts public speech into semi-private circulation while preserving the user's feeling that the post is public. The speaker may never know whether the post failed because it was uninteresting, because the audience was small, because the algorithm was indifferent, because foreign viewers could not see it, or because the platform quietly assigned it to a low-distribution category.

RT and PressTV reposting presents a vivid case because the author's own experience can be misleading. A United States–based user may post foreign state-media content every day without warnings or removals. From the user's side, this looks like evidence that the platform is not

interfering. But the absence of visible punishment does not prove the absence of containment. The posts may simply remain inside a stable domestic circle. They may be blocked or reduced in regions where the source is legally sensitive. They may be invisible to non-followers. They may not appear in recommendation systems. They may be permitted as personal expression while withheld from broader circulation.

Siloing is therefore not the opposite of censorship. It is censorship adapted to platforms that prefer deniability, scale, and risk management. It allows the platform to say that the user was not banned. It allows the user to believe he is still publishing. It allows governments to limit distribution without creating a dramatic censorship spectacle. It allows controversial material to exist without gaining public force. The post remains as a record, but its political energy is drained by confinement.

7.1. Defining the Silo

A silo is a bounded information chamber in which speech circulates among a limited population without reliably entering the wider public sphere. The boundary may be technical, legal, algorithmic, social, or behavioral. The defining feature is not absolute exclusion but constrained circulation. A siloed post is not necessarily hidden from everyone. It is hidden from enough people, in enough ways, that its apparent public status becomes misleading.

The silo differs from a private group. A private group announces its boundary. Members know they are speaking inside a limited space. A silo can occur inside a nominally public page. The post appears public, but distribution is restricted by systems the author cannot see. The page may have no explicit country restriction, no privacy limit, and no visible penalty. Yet the post may be effectively confined by ranking, recommendation, local-law restriction, source classification, or audience-history filtering.

The silo also differs from ordinary audience limitation. Every speaker has a finite audience. Not every low-reach post is censored. A post may fail to travel because it is dull, badly timed, poorly written, irrelevant, repetitive, or directed at a small topic. Siloing becomes a censorship problem when the limitation is not merely organic but structurally imposed by platform or legal systems that the user cannot audit. The difficulty is that organic limitation and imposed limitation can look identical from the author’s dashboard.

A useful definition must therefore include three elements. First, the content remains formally available to at least some viewers. Second, its broader circulation is limited by non-obvious systems. Third, the author lacks sufficient transparency to distinguish ordinary low reach from imposed containment. Siloing is not simply small audience size. It is small audience size produced or maintained by hidden distribution boundaries.

A silo may be built around the user, the source, the topic, the jurisdiction, or the audience. A user silo occurs when a page is generally downranked or confined to existing followers. A source silo occurs when material from RT, PressTV, or another disfavored outlet is allowed but not amplified. A topic silo occurs when political, war-related, sanctions-related, or foreign-policy content is generally reduced. A jurisdictional silo occurs when users in certain countries cannot see the post. An audience silo occurs when the post is shown primarily to people who already agree, already engage, or already belong to a platform-defined interest cluster.

These silos can overlap. A copied RT article may be source-siloed because of RT, jurisdictionally siloed because of European restrictions, algorithmically siloed because it is political content, socially siloed because only a stable audience engages with the page, and copyright-siloed because full copied text is treated as low-originality material. The author sees one post. The platform sees a bundle of risk signals.

The silo is powerful because it preserves plausible deniability. The platform can say the post was not removed. The government can say it did not jail the speaker. The user can say he has had no problem. All three statements may be technically true. Yet the post may still fail to reach the public. The central injury is not deletion but enclosure.

7.2. Geographic Siloing

Geographic siloing occurs when a post is visible in one location but restricted in another. This is the most literal form of legal geography. The platform does not treat the internet as one space. It divides visibility by country, region, legal regime, sanctions environment, and platform risk. A post may be available to the author in the United States but unavailable to a viewer in Germany, France, Poland, the United Kingdom, Canada, Australia, or another jurisdiction. The author's view remains normal because the restriction is not aimed at him.

This mechanism is especially relevant when source identity carries regional legal significance. RT is the strongest example. In parts of Europe, RT has been treated not merely as a controversial foreign news outlet but as part of a restricted Russian state-media apparatus. A platform that allows a U.S. user to post copied RT content may still decide that the same content should not be available to European viewers. The author may never know this unless a viewer in Europe checks the post directly.

PressTV can also be geographically siloed, though the pattern may differ. Its restrictions are more closely tied to sanctions, broadcast licensing, satellite carriage, domain disruption, and platform policy concerning Iranian state media. A copied PressTV article may not trigger the same geographic controls as RT, but it can still encounter regional caution in countries that have sanctioned or restricted Iranian state-linked media entities. The visibility layer may depend on the country, the source, the post format, and the platform's compliance posture.

Geographic siloing is difficult to detect because platforms personalize and localize content invisibly. The author cannot infer foreign visibility from domestic visibility. A viewer in the United States may see the post. A viewer in Canada may see it. A viewer in an EU country may not. Another EU viewer may see it if the platform has not detected the source. A logged-out viewer using a VPN may get a different result from a real local user with a normal account. A newly created foreign account may be treated as suspicious and therefore produce misleading evidence. The test environment matters.

Geographic siloing also complicates the meaning of public speech. A Facebook post may be public in a formal sense, but not public everywhere. The post may be globally addressable by a URL and still locally unavailable. It may exist in the platform database but be withheld in particular jurisdictions. It may appear in the author's archive while being filtered at the moment of viewing. Publicness becomes conditional.

This conditional publicness has political consequences. A writer may believe he is participating in an international debate, but his foreign audience may be partially absent. A post about NATO, Ukraine, Iran, Israel, sanctions, or media censorship may be visible mostly to domestic readers while unavailable or demoted abroad. This distorts feedback. The author may think European readers are uninterested when in fact they cannot see the post. Or he may think foreign restriction is occurring when the post simply lacks foreign audience demand. Without country-level analytics or external confirmation, the two possibilities are hard to separate.

Geographic siloing is therefore both a compliance tool and an epistemic trap. It allows platforms to obey local restrictions without global deletion. But it also prevents speakers from knowing the true geography of their audience. The post exists, but its map is hidden.

7.3. Algorithmic Siloing

Algorithmic siloing occurs when a platform's ranking and recommendation systems reduce a post's circulation without removing it. This is the most common and least visible form of soft censorship. Unlike geographic restriction, which may sometimes produce an unavailable-content notice, algorithmic restriction often produces nothing at all. The post simply travels poorly.

A social platform is not a chronological bulletin board. It is a prediction system. It estimates what users will engage with, what they may find meaningful, what they may report, what may cause harm, what may attract regulatory scrutiny, what may reduce user satisfaction, what may increase time on site, and what may create reputational risk for the company. Political content, state-media content, duplicate text, controversial foreign-policy material, and low-originality reposting may each receive ranking treatment that differs from ordinary personal content.

Algorithmic siloing can occur at several levels. A post may be shown to fewer followers. It may be excluded from non-follower recommendations. It may be prevented from appearing in suggested content modules. It may be ranked low in search. It may be excluded from trending surfaces. It may be shown only after higher-ranked content has exhausted user attention. It may not be eligible for boosting or advertising. It may be delayed in distribution. It may be deprioritized because it resembles spam or copied material. None of these actions requires deletion.

The author cannot easily observe these decisions. The dashboard may show reach, but not counterfactual reach. It may show how many people saw the post, but not how many would have seen it without demotion. It may show engagement, but not how many users were never offered the post. It may not explain why one post reached non-followers and another did not. It may not disclose whether state-media source identity, political topic classification, duplicate text, or prior account behavior affected ranking. The algorithmic border is therefore invisible by design or by complexity.

For RT and PressTV posts, algorithmic siloing may be triggered by several overlapping signals. The text may be recognized as copied from a state-linked outlet. The post may be classified as political or civic content. The full-article format may be classified as low-originality or duplicate content. The page's repeated behavior may look like systematic amplification. The audience may be politically homogeneous, leading the platform to keep distribution within that cluster. If users outside the cluster ignore or report similar posts, the system may learn not to distribute them broadly.

Algorithmic siloing can also be mistaken for social preference. A post that receives little engagement may have been shown to few people, or it may have been shown widely and ignored. The author cannot know without deeper metrics. Even when metrics exist, they may not reveal cause. This ambiguity is the core of soft censorship. The system does not need to tell the speaker "no." It can simply decline to carry the voice very far.

The platform can justify algorithmic siloing as quality control rather than censorship. It can say that users do not want political content, that duplicate content reduces feed quality, that state media requires caution, that civic content should be personalized, that harmful misinformation must not be amplified, or that recommendation systems are not obligated to promote every post. Some of these justifications may be reasonable. A platform has no duty to make every post equally viral. But when ranking choices systematically suppress disfavored political sources or controversial viewpoints without transparency, the boundary between curation and censorship becomes unstable.

Algorithmic siloing is therefore not simply a technical issue. It is a governance issue. Ranking is power. To rank is to distribute public attention. To demote is to alter political possibility. A post that is allowed but never carried is formally alive and practically weak. The speaker remains visible to himself, but not necessarily to the world.

7.4. Social Siloing and Self-Selected Audiences

Social siloing occurs when content circulates mainly among people who already expect, tolerate, or agree with it. This form of siloing can arise organically, algorithmically, or both. Social networks reward familiarity. Users engage with pages that match their interests. Algorithms learn from engagement and show future posts to similar users. Over time, a page develops a stable audience. The author sees regular reactions and assumes the post is reaching the public. In reality, it may be reaching a self-selected circle.

This is not always censorship. Communities naturally form around shared interests. A page devoted to foreign policy will not attract every reader. A page that posts RT and PressTV material will likely attract people already skeptical of mainstream Western narratives or interested in adversarial media. A small, stable audience can be the natural result of niche content. But social siloing becomes politically important when it combines with platform containment. The algorithm may interpret engagement from a narrow circle as evidence that the post should remain inside that circle. The audience's self-selection becomes the platform's distribution boundary.

Social siloing is psychologically powerful because it produces reassurance. The author receives likes. Familiar names comment. Some readers appreciate the posts. No one reports trouble. The page feels alive. But the same signs can indicate enclosure. The audience may be loyal but limited. The post may not be entering broader debate. It may not be challenging outsiders. It may not be reaching undecided readers. It may be circulating as confirmation within a known network.

This is why the statement "I have had no problem" must be interpreted carefully. A user may have no problem precisely because the posts are not escaping the expected audience. The content may not be threatening enough to the platform because it is contained. A post that reaches only those who already seek such material creates less public risk than the same post reaching millions of undecided viewers. Platforms may tolerate controversial speech inside small silos while suppressing its expansion.

Social siloing can also reduce reporting. If the post is shown mainly to sympathetic or tolerant users, fewer people will report it. If it reaches hostile audiences, reports may increase. Thus a lack of enforcement may reflect audience composition, not platform approval. The post remains

safe because it remains socially enclosed. The boundary is maintained by the same system that selects likely engagers.

For RT and PressTV content, social siloing has an additional effect: it may convert foreign state-media narratives into identity markers. Readers do not encounter the article simply as news. They encounter it as part of a shared posture: skepticism toward Western media, distrust of official narratives, interest in forbidden sources, or opposition to perceived censorship. The act of reading or sharing becomes a sign of belonging. The silo then protects not only content but identity.

This can happen across the political spectrum. Mainstream audiences can be siloed. Dissident audiences can be siloed. Pro-government and anti-government audiences can each inhabit curated information chambers. The issue is not which side is correct. The issue is that platform distribution increasingly organizes people into interpretive communities where content confirms the boundary of the group. In such a system, censorship may not need to suppress speech directly. It can allow each group to speak mainly to itself.

Social siloing is therefore both a comfort and a cage. It gives the speaker an audience. It also prevents the speaker from knowing whether he has reached beyond it. A page may feel public because people respond. But the real test is whether strangers, critics, foreign viewers, and non-followers encounter the material. Without that outward movement, the public sphere becomes a set of adjacent rooms whose occupants mistake room noise for public debate.

7.5. The Difference Between Being Blocked and Being Contained

Being blocked and being contained are different experiences. Blocking is a wall. Containment is a narrowing corridor. Blocking tells the user that access has failed. Containment lets the user continue while reducing the consequences of continuing. Blocking is visible, dramatic, and often contestable. Containment is quiet, ambiguous, and difficult to prove.

A blocked post is removed or inaccessible. A contained post remains available but with limited reach. A blocked source produces a clear question: who blocked it, under what authority, and why? A contained source produces a murkier question: why does this never travel, why do foreign viewers not respond, why is non-follower reach low, why do some people see it and others not, why does the dashboard lack the necessary data? Containment disperses responsibility across law, policy, ranking, audience behavior, and technical opacity.

The old censorship debate was organized around blocking because blocking was legible. It fit legal categories. It produced plaintiffs, defendants, orders, notices, bans, and public controversy. Containment is harder to litigate and harder to narrate. The platform can say it did not remove the speech. The government can say it did not ban the speaker. The user can still post. Critics

can say the content is simply unpopular. The harm lies in a missing counterfactual: the audience that might have existed but was never reached.

This is why containment is well suited to modern platform power. Platforms do not have to decide whether a controversial user may speak. They decide how far the speech travels. They do not have to ban all RT or PressTV-derived text. They can block official accounts, demote links, restrict regions, reduce recommendations, and leave small private reposts alone. They do not have to confront the user directly. They can place the user inside a low-risk distribution category.

Containment can be rational from the platform's perspective. A global platform must manage legal exposure, propaganda concerns, user trust, advertiser pressure, government demands, safety risks, and public criticism. Total removal creates backlash. Total openness creates regulatory risk. Containment offers a middle path. It allows the platform to preserve the appearance of openness while reducing the spread of material it considers risky. It is governance by throttling.

For the speaker, however, containment creates epistemic uncertainty. The speaker cannot know whether he is succeeding, failing, or being limited. He may continue posting daily for a year without trouble. That fact is meaningful but incomplete. It proves continuity of authorship, not freedom of distribution. It proves the post can exist, not that it can travel. It proves the room is open, not that the doors are.

A practical distinction follows. The question "Am I blocked?" asks whether the wall has appeared. The question "Am I contained?" asks whether invisible boundaries shape the audience. The first can often be tested by loading a page. The second requires comparison, outside verification, country checks, non-follower reach analysis, and post-type comparison. A VPN may test one layer. A real foreign viewer may test another. Dashboard data may show some patterns. None gives perfect knowledge. Containment is harder to measure because it is built from partial signals.

In the case of RT and PressTV reposting on Facebook, containment is the more plausible concern than complete blocking. A U.S.-based user may not be banned. The posts may remain visible domestically. Existing followers may continue to engage. But the posts may be withheld from some countries, demoted as state-media content, treated as low-originality copying, excluded from recommendation, or confined to a self-selected audience. This is not a contradiction. It is the defining structure of soft censorship.

The future of internet restriction is therefore not only the forbidden website. It is the visible post with an invisible ceiling. It is the public page that reaches only a private public. It is the

author who can still speak but cannot tell how far the sound travels. Containment does not silence the speaker. It reduces the world.

8. The Psychology of the Invisible Wall

The invisible wall is more difficult to confront than the visible ban because it leaves ordinary experience intact. The account opens. The page loads. The post publishes. The author can see it. Familiar readers respond. Nothing appears broken. No formal notice arrives. There is no obvious confrontation with a censor, regulator, or platform authority. The speaker continues to speak and therefore has no simple reason to believe that anything has changed. The wall is not encountered as a wall. It is encountered as normality.

This is the psychological genius of soft censorship. It does not create martyrs easily. It does not give the speaker a clear injury. It does not produce a blocked-page screenshot, a removal notice, or a government order that can be circulated as evidence. Instead, it produces uncertainty. The author wonders whether the post is not traveling because it is uninteresting, repetitive, unpopular, poorly timed, politically unfashionable, algorithmically disfavored, legally restricted, geographically unavailable, or quietly contained. The absence of evidence becomes part of the structure. The user cannot prove restriction because the platform does not disclose enough about distribution. He cannot prove freedom either, because his own view is not the public view.

The psychology of the invisible wall is therefore a psychology of inference. The user must interpret weak signals: low reach, stable audience, missing foreign engagement, absent country data, lack of warnings, lack of removals, repeated survival of posts, and occasional confirmation from others. None of these signals is decisive alone. Together they form a pattern, but the pattern remains ambiguous. A small audience may indicate censorship, or it may indicate limited demand. A post that survives may indicate tolerance, or it may indicate containment. A dashboard that lacks country data may indicate privacy thresholds, or it may hide the very evidence needed to diagnose geographic restriction.

This uncertainty affects behavior. Some users become more cautious, reducing what they post even without being told to stop. Others become more defiant, interpreting low reach as proof of suppression. Others normalize the situation and continue posting into a shrinking circle. Still others develop a ritual of publication detached from measurable audience: the act of posting becomes testimony, archive, or self-confirmation rather than communication. In each case, the invisible wall changes the speaker’s relationship to the public. The public becomes less knowable.

For RT and PressTV reposting, this psychology is especially strong because the user's lived evidence cuts in both directions. Daily posting without punishment suggests that the content is permitted. But the possibility of being siloed suggests that permission may not equal reach. The user is not asking whether he can speak. He already can. He is asking whether anyone outside the permitted room can hear him.

8.1. "I Have Had No Problem" as a Data Point

The statement "I have had no problem" is important and should not be dismissed. It is real evidence. It means the user has not experienced obvious deletion, suspension, warning, lockout, or direct platform sanction. It means the posts have been accepted at the point of publication. It means the account has not been treated in a visibly punitive way. If the user has posted similar material every day for a year, that is a significant record of non-removal.

But it is only one kind of evidence. It proves survival, not distribution. It proves that the posts can exist on the page, not that they travel beyond it. It proves that the platform has not chosen visible enforcement, not that invisible ranking systems treat the posts normally. It proves that the author has not been stopped, not that foreign viewers can see the posts. It proves continuity of posting, not completeness of audience.

This distinction is central. A post can survive because it is allowed. It can also survive because it is small. It can survive because it has not been reported. It can survive because copied text is harder to classify than a direct link. It can survive because the platform prefers demotion to removal. It can survive because it is visible only to people unlikely to object. It can survive because the enforcement system is inconsistent. Survival is evidence of one layer, but not of all layers.

The user's observation should therefore be treated as a starting point, not a conclusion. It narrows the question. The question is probably not, "Why is Facebook deleting these posts?" It is more likely, "How far do these posts travel?" or "Are these posts shown differently from ordinary posts?" or "Are they visible to viewers in countries where RT or PressTV faces restrictions?" The absence of obvious trouble shifts the investigation from hard censorship to soft containment.

There is also a psychological temptation to overvalue one's own experience because it is immediate. The author sees the post and therefore feels the post is real. A missing foreign audience is abstract. A hidden non-follower distribution penalty is abstract. A regional restriction is abstract unless someone reports it. Human beings naturally trust the visible interface over invisible counterfactuals. The platform benefits from that tendency. It gives the user enough visible success to sustain the feeling of publication.

At the same time, suspicion can run too far. Low reach is not always suppression. A post may fail because few people care, because the topic is repetitive, because the writing is long, because the audience is fatigued, because Facebook reduces political content generally, or because the post is not optimized for engagement. The phrase “I have had no problem” should not be reversed into “therefore I am definitely being censored.” It should be refined into “I have no evidence of overt enforcement, but I lack evidence of full visibility.”

This is the disciplined position. The user’s experience is valid. It shows the account is not openly blocked. It does not settle the distribution question. In a platform environment, the absence of trouble is compatible with freedom, neglect, tolerance, demotion, or containment. The task is to distinguish among them with better evidence.

8.2. Why the Author’s View Is Not the Audience’s View

The author’s view is privileged but misleading. It is privileged because the platform is designed to show the author his own content. It is misleading because that view is not the same as what other users see. On Facebook, the author’s page is not a neutral public window. It is a personalized interface generated for the account holder. It confirms that the post exists in the author’s environment. It does not confirm that the post appears in other environments.

Other viewers may encounter different versions of the same platform. A follower may see the post in a feed. A non-follower may never see it. A logged-out viewer may see only part of the page. A foreign viewer may see an unavailable-content notice. A viewer in a restricted jurisdiction may not see the post at all. A person searching for the page may find the page but not the post. A person clicking a direct link may see the post, while a person relying on feed distribution may not. A person using a mobile app may see different results from a person using a desktop browser. The post is not a single public object in practice; it is a set of conditional renderings.

This difference is difficult to grasp because social platforms simulate a public page. The interface invites the author to imagine that the post occupies one stable location. But behind the interface, the platform performs audience selection. It decides which users receive the post automatically, which users must seek it out, which users cannot see it, which users are shown it lower in the feed, which users are warned, and which users are never exposed. The author sees the object; the audience encounters distribution decisions.

This is why direct external checks matter. A real viewer in the United Kingdom, France, Germany, Canada, or Australia may see something different from the author. A non-follower may see something different from a follower. A person outside the author’s normal ideological network may see nothing unless given a direct link. A public post may be technically visible but not delivered. The author’s view cannot detect these distinctions.

The problem is intensified by missing analytics. If the dashboard does not show country-level reach, non-follower reach, or distribution explanations, the author cannot map the audience. He may know that some people engaged, but not who did not receive the post. He may know that a post reached a certain number of viewers, but not whether it was withheld from regions. He may know that a familiar group responded, but not whether strangers were excluded. The most important audience data is the audience that never appears.

This creates a philosophical problem of negative evidence. The absent viewer leaves no trace. A German reader who never saw the post cannot comment that he did not see it. A non-follower who was never offered the post cannot report non-exposure. A search result that never displays the post generates no notification. The platform records such events internally, but the user does not see them. The author's world is built from positive traces: likes, comments, visible posts. Containment operates through missing traces.

The author's view is therefore psychologically comforting but epistemically weak. It confirms that the post exists for the author. It does not confirm that the post exists as a public object with equal availability across audiences. The responsible conclusion is not paranoia, but humility: on a platform, one never sees the whole audience architecture from the author's side.

8.3. The Comfort of Apparent Normalcy

Apparent normalcy is powerful because it allows both the platform and the user to avoid confrontation. The platform avoids accusations of censorship by leaving the post visible. The user avoids the anxiety of direct suppression because daily practice continues. The system appears stable. The author posts. The page updates. The familiar audience reacts. The pattern repeats. Over time, this repetition becomes evidence in the user's mind that nothing is wrong.

This comfort is not irrational. Human beings rely on routine to judge systems. If a door opens every day, one assumes the door is not locked. If posts publish every day, one assumes publication is real. If no warning appears, one assumes no rule has been violated. If no one complains, one assumes the audience can see the material. These assumptions are reasonable in ordinary physical settings. They are less reliable on algorithmic platforms.

The platform interface is designed to preserve continuity. It does not normally announce every ranking decision. It does not tell the user, "This post was shown to fewer non-followers because it contains duplicated political text." It does not say, "This post may be unavailable in certain countries." It does not say, "This source category is sensitive." It does not say, "Your page is allowed to publish but unlikely to be recommended." The user receives a smooth interface covering a fragmented distribution system.

Apparent normalcy can therefore become a form of anesthesia. It reduces the emotional signal that something may be wrong. A hard block produces anger, fear, or resistance. Soft

containment produces adaptation. The user may unconsciously lower expectations. He may come to accept a small circle as the natural audience. He may stop asking whether foreign readers see the posts. He may treat the absence of growth as ordinary platform decline. He may continue posting as ritual rather than as public communication.

This has political consequences. If speech is formally permitted but practically enclosed, dissent can be domesticated without being crushed. The speaker remains active, but the speech loses expansion power. The platform can claim openness. The state can avoid the spectacle of repression. The audience sees little conflict. The post exists as a private-public artifact: public enough to satisfy the speaker, private enough to reduce impact.

The comfort of normalcy also creates conflict between intuition and measurement. The user may feel that everything is working because his daily routine works. A measurement-oriented approach asks a colder question: compared with what? Do RT-derived posts perform differently from ordinary posts? Do PressTV-derived posts travel less far than personal commentary? Do full copied articles reach fewer non-followers than short excerpts? Can foreign contacts see the posts? Are shares possible abroad? Do posts appear in search? Without comparison, normalcy is only a feeling.

The comfort of normalcy is especially likely when the audience is stable. Familiar readers create social proof. A small but loyal group can make a post feel public. In reality, the group may be the boundary. The same names appear because the platform has learned that those names engage. Their engagement sustains the post inside the silo. The author sees life; the system sees containment.

The goal is not to replace comfort with paranoia. The goal is to replace comfort with measurement. Apparent normalcy should be respected as lived experience but tested as a distribution claim. The post may indeed be free. It may also be contained. Only structured comparison can begin to tell the difference.

8.4. Containment Without Notification

Containment without notification is the purest form of soft censorship. The user is not told because the platform has not formally removed anything. The content may remain technically available. The account may remain active. The platform may classify the restriction as ranking, quality control, local compliance, recommendation eligibility, civic-content reduction, duplicate-content handling, or user-preference optimization. None of these necessarily produces a notice.

Notification is usually tied to visible enforcement. A post is removed; the user receives a notice. An account is suspended; the user receives an explanation. A copyright complaint is filed; the user receives a claim. A legal restriction applies; the platform may disclose that content is unavailable in a region. But ranking decisions are often treated as internal product decisions.

The platform does not notify every user whose post is shown to fewer people. It does not explain every distribution choice. If containment is implemented through ranking, it may leave no message.

This lack of notification changes accountability. A removed post can be challenged. A demoted post is harder to challenge because the user may not know demotion occurred. A country-specific block may be discoverable if a foreign viewer reports it, but otherwise remain invisible. A recommendation exclusion may be impossible to prove externally. A reduced distribution signal may be hidden inside aggregate metrics. The user cannot appeal what he cannot identify.

Containment without notification can arise for many reasons. The post may include copied material. It may be political. It may originate from a state-linked source. It may be associated with a restricted entity. It may trigger low-originality classification. It may be reported by some users but not removed. It may be held back from recommendations because the platform predicts negative feedback. It may be limited in foreign jurisdictions. It may be affected by account-level trust scores. Each reason has a different policy logic, but the user experiences the same silence.

This silence can produce two opposite psychological errors. The first is complacency: because no notice arrived, nothing happened. The second is overinterpretation: because reach is low, censorship must be happening. Both errors come from the same opacity. The platform does not provide enough information for confident judgment. The rational response is neither complacency nor paranoia, but experimental testing.

For the RT and PressTV case, containment without notification is plausible because it fits the user's observations. Daily posts remain visible. There are no warnings. The dashboard lacks country data. The user suspects siloing but cannot verify it. This is exactly the kind of situation soft censorship creates: enough visibility to continue, not enough transparency to know the audience.

The moral question is whether platforms owe users notice when distribution is intentionally reduced for political, legal, or source-based reasons. One can argue that platforms have no obligation to promote every post. They must rank content somehow. Not every ranking decision can generate a notice. But when distribution is reduced because of source identity, state-media classification, legal geography, or political sensitivity, the decision begins to resemble governance rather than mere product design. In such cases, lack of notification weakens democratic accountability.

Containment without notification is therefore not only a technical issue. It is a civic issue. Public discourse increasingly occurs inside systems that can alter reach invisibly. The speaker may be

present. The post may be present. The public may be absent. The platform may know why. The user may never know.

This is the invisible wall in its mature form: not a barrier that stops the speaker, but a boundary that prevents the speaker from knowing whether speech has become public. It leaves no dramatic wound. It produces no obvious scandal. It permits the daily ritual of posting. And because it permits the ritual, it can be mistaken for freedom.

9. Testing Visibility Without Creating False Signals

Testing platform visibility is difficult because the object being tested is not simply whether a post exists. Existence is easy to confirm. The author can see the post. A friend can click it. A direct link may open. The harder question is whether the post is distributed normally, visible across jurisdictions, available to non-followers, eligible for recommendation, and treated like comparable posts that do not contain RT or PressTV material. A poor test can create false reassurance or false alarm. The design of the test matters.

The first principle is that visibility must be tested from outside the author's perspective. The author's account is the least useful viewpoint because it is privileged. The platform is designed to show users their own posts. A successful author-side view confirms publication but not reach. To evaluate siloing, the test must include viewers who differ from the author by geography, account history, follower status, device, login state, language, and social relationship to the page. A post that appears to the author may not appear to a foreign viewer. A post that opens through a direct link may not appear in feeds. A post that is visible to followers may not be visible to non-followers.

The second principle is comparison. One RT post proves little. One PressTV post proves little. One ordinary post proves little. Visibility must be compared across categories. The user should compare ordinary personal posts, original commentary posts, copied RT articles, copied PressTV articles, short excerpts with commentary, and full verbatim reposts. If all posts perform similarly, siloing is less likely. If RT or PressTV posts consistently show reduced reach, fewer shares, weaker non-follower exposure, or foreign invisibility, the case for containment becomes stronger. The test should not ask merely whether one post can be seen. It should ask whether different kinds of posts behave differently.

The third principle is avoiding artificial signals. A VPN, a new foreign account, or a logged-out browser may test some technical layers, but they do not fully reproduce a normal foreign user. Facebook uses more than IP address. It may consider account age, region history, device identity, language, login pattern, friend network, page interaction history, and suspicious-account signals. A new account created through a VPN may be treated as abnormal. A logged-

out viewer may see less content than a logged-in viewer. A VPN may show whether blunt geoblocking exists, but it may not reveal feed distribution or real audience reach.

The fourth principle is documenting outcomes precisely. The test should not rely on impressions such as “it seems visible” or “it looks suppressed.” It should record whether the post opens from a direct link, whether it appears on the public page, whether it appears in search, whether the share button is available, whether a viewer sees a warning, whether the viewer can comment, whether the post appears in feed without a direct link, and whether the viewer is logged in or logged out. Each of these observations tests a different layer.

The goal is not to prove censorship from one anomaly. The goal is to build a pattern. Platform siloing is probabilistic, partial, and contextual. It may not appear every time. A structured test cannot reveal everything the platform does internally, but it can reduce uncertainty. It can show whether the posts are merely published or actually traveling.

9.1. Why a VPN Is an Imperfect Test

A VPN is useful but limited. It can make a user’s internet traffic appear to originate from another country. If a post or website is blocked by crude IP-based geolocation, a VPN may reveal that difference. A user in Kansas can connect through a United Kingdom, German, French, Canadian, or Australian endpoint and test whether a public post or webpage loads differently. This can detect some forms of geographic blocking.

But Facebook is not a simple website that uses only IP address to decide what a user sees. It is an account-based platform. When a user logs in, the platform already knows the account’s history, normal location, language, device, behavior, page relationships, and social graph. A U.S. account suddenly appearing through a UK VPN does not become a normal UK account. It becomes a U.S. account using an unusual connection. Facebook may still apply the user’s account region, not merely the VPN endpoint. It may also treat sudden location shifts as suspicious, which can distort the result.

Logged-out VPN testing has its own weakness. Many Facebook pages and posts are less visible to logged-out users than to logged-in users. A failure to see a post while logged out through a VPN may reflect Facebook’s general restrictions on logged-out browsing rather than censorship of the post. The post may be visible to a real local user but not to a logged-out browser. Conversely, a post that opens logged out may still not be delivered in feeds to local users. Loading a direct link is not the same as distribution.

A VPN also cannot test algorithmic reach. It may show whether a direct post URL is available from a country. It cannot show whether Facebook would recommend that post to local users, show it in feeds, allow it to spread through shares, or include it in search. A post can pass a VPN direct-link test and still be algorithmically contained. The VPN tests access, not circulation.

A VPN may also create false confidence. If the post opens through a UK endpoint, the user might conclude that UK viewers can see it normally. But a direct URL opened through a VPN is not the same as normal discovery. A real viewer may not see the post in feed. The post may be visible only if the viewer already has the link. That is not full public visibility. It is addressable access.

The VPN test is therefore best treated as a low-level diagnostic. It can answer narrow questions: does the direct link load from this apparent country, and does the public page display the post under these conditions? It cannot answer the broader questions: does the post reach non-followers, is it recommended, is it demoted, is it visible to ordinary local accounts, and does it circulate normally? The VPN is a flashlight, not a map.

9.2. Why a New Foreign Account Can Mislead the Experiment

Creating a new foreign Facebook account to test visibility may seem logical, but it can produce misleading results. A new account is not a normal viewer. It has no history, no stable social graph, no ordinary engagement pattern, no established location, and no trust record. If created through a VPN, it may also look suspicious. Platforms are designed to detect fake, duplicate, spam, or manipulation accounts. A newly created foreign account may be restricted, challenged, limited, or shown an abnormal version of Facebook.

Such an account may fail to see posts for reasons unrelated to the post itself. It may lack permission to view certain public material because Facebook limits new accounts. It may be prevented from searching effectively. It may not be shown political content because it has no interaction history. It may trigger security checks. It may be classified as low-trust. If the account cannot see the post, the user might wrongly conclude that the post is geoblocked. If it can see the post, the user might wrongly conclude that all normal foreign users can see it. Both conclusions would be weak.

A new account also does not test feed distribution. It has no normal feed. It does not follow the page organically. It has not developed interests. It does not represent a real UK, EU, Canadian, or Australian viewer. It represents an artificial identity created for testing. Platform systems may treat such identities differently from ordinary users. The result may say more about account-integrity systems than about RT or PressTV visibility.

There is also a policy problem. Platforms generally discourage or prohibit users from maintaining multiple personal accounts, especially accounts that misrepresent identity or location. Even if the purpose is harmless research, the account itself may violate platform rules. A test that creates rule-breaking behavior can contaminate the result. If the account is restricted, the restriction may be caused by the test method rather than by the content being tested.

A better method is to use real viewers. A real person in the United Kingdom, Germany, France, Canada, Australia, or another target country can test the post from a normal account, normal device, normal location, and normal social context. That person's experience is much closer to the audience being studied. If several real viewers in different jurisdictions report the same result, the evidence becomes stronger.

The test should also include viewers with different relationships to the page. A follower may see more than a non-follower. A friend may see more than a stranger. A person who has previously interacted with the page may see more than someone encountering it for the first time. Real-viewer testing should therefore specify whether the person follows the page, has interacted with it before, found the post by direct link, found it by visiting the page, or saw it naturally in feed.

A new foreign account seems controlled, but it is actually noisy. A real viewer seems less controlled, but is more ecologically valid. For testing platform siloing, ecological validity matters. The question is not whether an artificial test account can open a post. The question is whether ordinary people in the relevant countries can see and share it under normal conditions.

9.3. Human Verification by Real Viewers in Target Countries

Human verification is the strongest practical test available to an ordinary user. It does not reveal the platform's internal ranking system, but it can detect real-world visibility differences that the author cannot observe. The key is to ask specific people in specific countries to perform specific actions and report precise results.

The test should begin with a fixed set of posts. The user should choose several ordinary posts, several original-commentary posts, several copied RT posts, and several copied PressTV posts. Each post should have a date, title or topic, direct link, and classification. The viewer should not be asked vaguely whether the page looks normal. The viewer should be asked to test each selected post in the same way.

A good human test has several steps. First, the viewer opens the page normally while logged into Facebook. Second, the viewer checks whether the selected post appears by scrolling the page, not by using the direct link. Third, the viewer opens the direct post link and reports whether it loads. Fourth, the viewer reports whether the post can be liked, commented on, and shared. Fifth, the viewer checks whether any warning, unavailable message, missing-content notice, or unusual label appears. Sixth, the viewer records whether they follow the page, have interacted with it before, and what country they are in.

The distinction between page browsing and direct-link access is important. A post may open through a direct link but not appear naturally on the page or in feed. Direct-link access shows that the object exists for that viewer. Page browsing shows that the object is publicly displayed

in context. Feed appearance shows distribution. Search appearance shows discoverability. These are separate layers.

The test should include multiple countries where restrictions differ. For RT, EU viewers are especially important because European restrictions have been more extensive. The United Kingdom and Canada should be tested separately because their regulatory histories differ from the EU. Australia is useful for PressTV and sanctions-related visibility questions. A U.S. control viewer is also necessary. If U.S. viewers can see the post and EU viewers cannot, geographic siloing becomes plausible. If all viewers can see it through direct links but only followers see it in feeds, algorithmic or social containment becomes more plausible.

The test should include more than one person per region if possible. One viewer's result may reflect account settings, privacy limits, device issues, language settings, or temporary glitches. Three viewers in the same region reporting the same result provide stronger evidence. A mix of followers and non-followers is also useful. If followers can see the post but non-followers cannot find it, the post may be socially or algorithmically contained rather than legally blocked.

Human verification should be repeated over time. A post may be visible when first published and restricted later. A legal request may arrive after publication. An automated system may classify the post later. A copyright claim may remove it later. A platform update may change distribution. Therefore testing one time only provides a snapshot. Testing at publication, after twenty-four hours, after one week, and after one month can reveal delayed restriction.

The results should be recorded in a simple table. The table should list post type, source, date, country of viewer, follower status, direct-link visibility, page visibility, feed visibility, share availability, warning message, and comments. The goal is to replace impression with evidence. Even if the evidence remains incomplete, it will be better than relying on the author's view.

Human verification has limits. It cannot prove why a post is missing. It cannot reveal internal demotion. It cannot prove intent. But it can answer the most practical question: can real people in the target countries see this post under normal conditions? For a user concerned about siloing, that is the first question that matters.

9.4. Logged-Out Checks, Direct Links, and Shareability Tests

Logged-out checks test public availability without relying on a particular Facebook account. They can be useful, but they must be interpreted carefully. Facebook often limits what logged-out users can see. It may show only part of a page, block scrolling, require login, hide comments, or refuse to display certain posts. A logged-out failure may not mean the post is restricted because of RT or PressTV. It may mean Facebook does not expose that material to anonymous visitors.

Still, logged-out checks can reveal blunt differences. If a post opens while logged out from a U.S. connection but not from a UK or EU VPN endpoint, geographic restriction becomes possible. If ordinary posts open while RT-derived posts do not, source-specific restriction becomes more plausible. If no posts open logged out from any location, the test says little about censorship and more about Facebook's general access design.

Direct-link tests are stronger than page-level browsing for confirming whether a particular post exists for a viewer. The author sends the exact post link. The viewer opens it and reports the result. If the link opens, the post is addressable for that viewer. If the link produces an unavailable message, the post may be restricted, deleted, privacy-limited, or otherwise inaccessible. If the link opens for U.S. viewers but not for EU viewers, geographic siloing becomes plausible.

But direct-link visibility is not full visibility. A post can open by direct link while being absent from feeds, search, recommendations, and page timelines. That means it is technically available but practically weak. Direct-link access is like a book stored in a locked archive where the key works only if someone already knows the shelf number. It is not the same as public circulation.

Shareability is another important layer. A viewer may be able to read a post but not share it. Or the share may appear to work but produce limited distribution. The share button itself can reveal platform treatment. If ordinary posts can be shared but RT-derived posts cannot, that suggests source or policy restriction. If a viewer can share the post but the shared version is visible only to a narrow audience, a second layer of containment may appear. Testing shareability requires care because privacy settings, viewer account type, and page configuration can also affect results.

A useful shareability test asks the foreign viewer to attempt three actions: share the post publicly, send it by Messenger, and copy the link into a new post with neutral commentary. The viewer should not actually publish anything that could create risk unless willing to do so. The point is to observe whether Facebook blocks or warns at the moment of sharing. If a warning appears, it should be copied exactly. If no warning appears, that should also be recorded.

Search testing can add another layer. The viewer can search the page name plus a distinctive phrase from the post. If ordinary posts appear but copied RT or PressTV posts do not, discoverability may differ. Search is unstable and personalized, so the result is not conclusive. But repeated differences across viewers can be meaningful.

The strongest practical method combines logged-in real-viewer checks, direct-link tests, page browsing, shareability tests, and limited logged-out VPN checks. Each test sees one layer. Together they create a profile. A post may be visible by direct link but not on the page. It may be

visible to followers but not non-followers. It may be visible in the United States but not the EU. It may be readable but not shareable. Each outcome corresponds to a different kind of silo.

The central discipline is not to overinterpret any single result. A failed logged-out check is not proof of censorship. A successful direct-link check is not proof of normal distribution. A share button does not prove reach. A missing search result does not prove suppression. Siloing is inferred from patterns across tests.

9.5. Comparing Ordinary Posts, RT Posts, and PressTV Posts

Comparison is the heart of the experiment. Without comparison, the user cannot know whether RT and PressTV posts are being treated differently from the rest of the page. A post may have low reach because the page is small, the topic is narrow, the timing is poor, or the audience is tired. Only comparison can separate general page behavior from source-specific effects.

The simplest design uses three categories: ordinary posts, RT-derived posts, and PressTV-derived posts. A better design uses five categories: ordinary personal or general posts, original foreign-policy commentary without copied articles, short RT excerpts with commentary, full RT reposts, short PressTV excerpts with commentary, and full PressTV reposts. This allows the user to test not only source identity but form. If full reposts perform worse than excerpts with commentary, the problem may be copied-article format. If RT performs worse than PressTV, the problem may be source-specific. If all political posts perform worse than ordinary posts, the problem may be civic-content ranking. If all posts perform similarly, visible containment is less evident.

Each category should include enough posts to reduce randomness. Ten posts per category is a reasonable starting point. More is better. The user should record date, time, post length, source, topic, whether the post contains a URL, whether it contains copied full text, whether it contains original commentary, reach, reactions, comments, shares, follower versus non-follower exposure if available, and any visible warnings. If country data is unavailable, external viewer checks can substitute for part of that missing information.

The comparison should control for timing as much as possible. Posts made at different times of day or on different days may perform differently. A fair test should compare posts made under similar conditions or average across enough posts that timing effects matter less. Topic intensity also matters. A dramatic war headline may outperform a routine policy article regardless of source. The user should note major news events that could affect engagement.

One useful metric is share ratio. If ordinary posts are shared but RT-derived full reposts are rarely shared, that may indicate either audience hesitation or platform friction. Another useful metric is non-follower reach. If ordinary posts reach non-followers but copied state-media posts

remain among followers, that suggests containment. A third useful metric is external visibility. If real EU viewers can see ordinary posts but cannot see RT-derived posts, geographic siloing becomes more plausible.

The user should also compare full reposts with rewritten summaries. A rewritten summary in the user's own words, with source attribution and analysis, may travel differently from verbatim copying. If summaries perform better than full copies, this supports the idea that platform systems penalize duplicate or mirror-like content. It also supports the practical recommendation to shift from republication to commentary.

PressTV and RT should not be assumed to behave identically. RT may trigger stronger European legal sensitivity because of the Ukraine-war restrictions. PressTV may trigger different sanctions, state-media, or reputational signals. If PressTV posts travel more freely than RT posts among European viewers, that distinction matters. If both behave similarly, the issue may be broader state-media or copied-content handling.

The experiment should include a negative control. The user can copy full articles from an ordinary nonrestricted source, such as a mainstream news outlet, and compare performance. If all full copied articles perform poorly, the issue may be duplicate-content or copyright risk rather than state-media restriction. If only RT or PressTV copies perform poorly, source identity becomes more likely. If mainstream copied articles trigger copyright complaints while RT or PressTV articles do not, that reveals a different enforcement pattern.

The final outcome should be a visibility profile, not a verdict. The profile might show that posts are visible domestically but not in the EU. It might show that they are visible by direct link but not distributed to non-followers. It might show that full copied articles underperform commentary. It might show no detectable difference. Each result would refine the theory.

Testing visibility is therefore an empirical discipline. It replaces the vague feeling of being blocked or free with structured observation. It does not require creating a foreign account. It does not require violating platform rules. It requires comparison, real viewers, careful records, and modest conclusions. The question is not whether the post exists. The question is how far it travels, where it stops, and whether the stop leaves a trace.

10. Measurement Limits Inside the Facebook Dashboard

The Facebook dashboard gives the user the impression of measurement, but not necessarily the measurement needed to diagnose siloing. It may show reactions, comments, shares, reach, impressions, follower growth, content interactions, and post performance. These are useful numbers, but they do not fully answer the key question: who did not see the post, and why?

The dashboard can confirm that a post received some circulation. It cannot always reveal whether the post was withheld from certain countries, excluded from recommendations, limited to followers, downranked because of source identity, suppressed as duplicate content, or treated differently from comparable posts.

This limitation is not accidental in the practical sense. Platforms measure far more internally than they disclose externally. Facebook must know enough to rank posts, serve ads, detect spam, enforce local law, manage copyright claims, classify content, estimate user interest, and satisfy regulatory reporting requirements. But the page administrator receives only a simplified surface. The platform does not expose the full logic of distribution. It does not show the counterfactual audience. It does not explain every demotion. It does not identify every legal or source-based risk signal. It does not tell the user how many people would have seen the post had the content been ordinary, original, domestic, nonpolitical, or non-state-media-derived.

For the RT and PressTV case, this gap is decisive. A page owner may see that a copied article was published and received some engagement. But the dashboard may not show whether European viewers were excluded, whether non-followers were reached, whether the post was recommended, whether the post was shareable in Canada, whether Australian viewers saw it differently, whether the platform recognized the source, or whether full-text copying reduced distribution. The dashboard can document positive exposure. It has difficulty documenting invisible absence.

A missing country breakdown is especially important. Without country-level data, the user cannot easily test geographic siloing from inside the platform. If a post receives 400 impressions, the number alone does not reveal whether those impressions came entirely from the United States or included viewers in the European Union, United Kingdom, Canada, Australia, or other relevant jurisdictions. If RT-derived posts never reach Europe, but the dashboard does not show country distribution, the user may never know. The absence of geographic data becomes part of the silo.

The dashboard is therefore not a neutral window into the public. It is a managed instrument. It gives enough information to evaluate engagement, but not necessarily enough to audit governance. For ordinary page management, that may be acceptable. For diagnosing soft censorship, it is inadequate. The user must supplement dashboard numbers with external tests, real-viewer verification, comparative post analysis, direct-link checks, and careful interpretation of patterns across time.

10.1. Missing Country-Level Data

Country-level data is essential for testing whether a post crosses legal and geographic boundaries. In the RT and PressTV case, the most important audiences are not only domestic

followers but viewers in jurisdictions where the outlets face restrictions, sanctions, broadcast limitations, or reputational treatment. If the dashboard does not show countries, the user cannot directly determine whether those viewers are present or absent. The post may appear successful while never leaving the United States.

This is not a minor technical inconvenience. Geographic visibility is one of the central questions of the paper. A post can be fully visible in Kansas and unavailable in Berlin. It can reach Missouri, Texas, and California while failing to appear in France, Poland, or the Netherlands. It can be visible to a U.S. follower traveling abroad but not to a normal local account. It can be accessible by direct link in Canada but not delivered to Canadian feeds. Without country data, these distinctions disappear into a single aggregate number.

Aggregate reach is therefore misleading. A post with 1,000 impressions may appear more public than a post with 100 impressions, but if all 1,000 impressions came from the same domestic audience, the post may still be geographically siloed. Conversely, a post with modest reach may be more internationally distributed if it reaches several countries. The dashboard's failure to show geography prevents the user from distinguishing domestic circulation from cross-border circulation.

The absence of country data can have several explanations. The page may not have enough audience volume to meet privacy thresholds. The dashboard view may be limited to post performance rather than audience demographics. The account may be using a profile, professional mode, or page type that does not expose the relevant metric. The interface may have changed. Meta may suppress granular demographics for privacy, safety, or policy reasons. Whatever the reason, the result is the same: the page owner cannot see the map.

This missing map affects interpretation of the user's daily posting experience. If the author posts RT or PressTV material every day and sees no problem, he may infer that the posts are generally available. But if the dashboard lacks country data, he cannot know whether the posts reach countries where restrictions matter. He is measuring survival, not geography. He is seeing domestic continuity, not global distribution.

The missing data also limits the ability to compare RT and PressTV. RT may face stronger restrictions in some European settings. PressTV may face different sanctions or platform treatment. If the dashboard showed countries, the user could compare whether RT-derived posts underperform in Europe relative to PressTV-derived posts. Without that data, the comparison must rely on external viewers or indirect reach patterns.

The practical conclusion is that missing country data shifts the burden to outside verification. The page owner should not assume international visibility from aggregate metrics. He should ask real viewers in target countries to check posts directly. He should compare ordinary posts

with RT-derived and PressTV-derived posts. He should record which countries can see direct links, page timeline entries, and share buttons. The dashboard's silence must be treated not as reassurance, but as an unknown.

10.2. Thresholds, Privacy Limits, and Page-Type Restrictions

Facebook does not necessarily reveal all demographic or geographic metrics to every page owner. Platforms often use thresholds to protect privacy. If too few people from a region interact with a page, showing country-level data could expose or imply information about identifiable individuals. A page with a small audience may therefore receive aggregate numbers while losing geographic breakdowns. The platform may be protecting user privacy, but the effect is also to reduce the page owner's ability to diagnose distribution boundaries.

Thresholds create a paradox. The smaller the page, the more likely it is to need detailed evidence to understand whether it is being contained. But the smaller the page, the less likely the dashboard is to provide granular data. A large page may see country distribution because enough viewers exist in each category. A small page may see only aggregate reach. Yet small pages are precisely where soft containment can be hardest to distinguish from ordinary obscurity. Low reach, missing geography, and stable audience can all look natural.

Privacy limits are not necessarily illegitimate. Revealing that one person in a small town or one person in a country viewed a politically sensitive post could create risk. Platforms must avoid exposing users through overly granular reporting. But privacy protection and platform opacity can produce the same outcome for the speaker. The page owner cannot tell whether foreign viewers are absent because no one saw the post, because the platform restricted the post, or because the platform refuses to display the data.

Page type also matters. A full business page, public page, professional-mode profile, community page, group, and personal profile may expose different analytics. Some interfaces show audience data. Others show only post performance. Some metrics appear in Meta Business Suite but not in a simplified dashboard. Some features depend on follower count, activity level, account status, or region. The user may believe he has "the Facebook dashboard," but there are multiple dashboards, each with different measurement power.

Interface changes worsen the problem. Platforms frequently redesign dashboards, rename metrics, move reports, remove exports, or change thresholds. A user may read advice about analytics that no longer matches the current interface. Country data may exist in one view and not another. Export options may appear for some pages but not others. A metric may be available only after a threshold is reached. The user experiences this as confusion, but the deeper issue is that the platform controls the measuring instrument.

These limits shape the interpretation of RT and PressTV posting. If a copied article receives modest engagement, the user cannot know from a restricted dashboard whether the post was limited to a domestic core audience or actually reached foreign viewers. If a post receives no foreign comments, the reason may be absence of foreign reach, lack of interest, language differences, or dashboard limitations. If PressTV posts appear similar to ordinary posts, that may reflect real equality or insufficient measurement resolution.

Thresholds also make controlled experiments harder. A user might compare ten ordinary posts and ten RT posts, but if the dashboard reports only aggregate reach and interactions, the experiment may detect only large differences. Subtle country-specific restriction may remain hidden. A post could be blocked in the European Union but still show similar aggregate reach if most of the page's audience is American. Without geography, the European absence disappears.

The solution is not to reject the dashboard, but to understand its limits. Dashboard data should be treated as partial evidence. It can show whether posts receive engagement. It can help compare relative performance. It can identify major reach differences. But it cannot, by itself, prove the absence of siloing. It is a surface instrument operating inside the very system being tested.

10.3. Reach, Engagement, Shares, and Non-Follower Exposure

The metrics most likely to help diagnose siloing are reach, engagement, shares, and non-follower exposure. Each measures a different layer of circulation. Reach measures how many accounts were exposed to the post. Engagement measures how many acted on it. Shares measure whether viewers helped move it outward. Non-follower exposure measures whether the post escaped the existing audience. None is perfect alone, but together they can reveal patterns.

Reach is the broadest metric, but it is also the most ambiguous. High reach suggests the post was distributed to many accounts. Low reach suggests limited distribution. But reach does not reveal why distribution was high or low. A post may have low reach because the platform demoted it, because followers ignored it, because the topic was narrow, because it was posted at a poor time, or because the page is small. Reach must be compared across similar posts. The question is not whether one RT post had low reach. The question is whether RT-derived posts consistently have lower reach than comparable ordinary posts.

Engagement adds another layer. A post with high reach and low engagement may have been shown widely but failed to interest viewers. A post with low reach and high engagement may be liked by a small core audience but not distributed outward. A post with modest reach and strong comments from the same names may indicate social siloing. Engagement quality also matters. Comments from followers suggest internal circulation. Comments from strangers

suggest expansion. Reactions alone can be misleading because a loyal audience may react habitually.

Shares are especially important because they are one of the mechanisms by which a post leaves the original audience. A post that receives likes but no shares may remain inside the page's existing network. A post that is shared by viewers can enter new networks. If RT or PressTV posts are rarely shared compared with ordinary posts, several interpretations are possible. Viewers may hesitate to share controversial state-media material. The platform may reduce share visibility. The post may be less engaging. The source may carry stigma. The article may be too long. Again, comparison is necessary.

Shareability should be distinguished from actual shares. A post may have a share button but still not travel. A viewer may share it, but the shared post may receive little distribution. The platform may allow sharing but demote the shared object. Therefore the user should not only count shares but, when possible, ask viewers whether the share process produces warnings or restrictions. If foreign viewers can read but not share, that is a different form of containment.

Non-follower exposure may be the most important metric. A post that reaches followers is circulating inside the existing silo. A post that reaches non-followers has at least some chance of entering the broader public. If the dashboard shows that ordinary posts reach non-followers but RT or PressTV full reposts do not, that is evidence of containment. If excerpts with commentary reach non-followers while full copied articles do not, that suggests form-based restriction. If all political posts have low non-follower reach, the cause may be general civic-content downranking rather than source-specific action.

The follower versus non-follower distinction also clarifies the user's experience. A stable audience can make posts feel successful even when they do not expand. If the same followers see the posts every day, the user receives feedback. But the feedback may come from inside the silo. Non-follower metrics reveal whether the post crosses the boundary. Without them, the user may mistake loyal internal circulation for public reach.

These metrics should be recorded over time. One post is too noisy. A thirty-post comparison can reveal patterns. The user should track ordinary posts, RT full reposts, RT excerpts with commentary, PressTV full reposts, PressTV excerpts with commentary, and original analysis. For each, record reach, reactions, comments, shares, non-follower exposure if available, and any external viewer results. The purpose is not to produce a perfect scientific study, but to replace intuition with structured evidence.

The metrics should also be interpreted modestly. Low non-follower reach does not prove censorship. High non-follower reach does not prove absence of restriction. A post can be partially restricted and still travel. A post can travel domestically while blocked abroad. A post

can get shares among followers but no recommendation expansion. Measurement inside a platform is always incomplete. But incomplete measurement is still better than relying on the author's own view.

10.4. The Problem of Measuring What the Platform Does Not Report

The deepest measurement problem is the counterfactual. The user wants to know not only how many people saw the post, but how many would have seen it if the post had not contained RT or PressTV material, had not been copied in full, had not been political, had not been associated with a restricted source, or had not been subject to regional legal concern. The platform does not report this counterfactual. It reports the outcome, not the suppressed possibility.

This is the core difficulty of proving soft censorship. Hard censorship leaves evidence: a removal notice, a blocked page, a seized domain, a legal order. Soft censorship leaves a missing audience. The absence of an audience is hard to measure because the missing viewers did not produce data visible to the user. The platform may know that it did not recommend the post to non-followers. It may know that it withheld the post in certain regions. It may know that duplicate-content systems reduced ranking. But the page owner may see only the final reach number.

The platform also does not usually reveal the causal chain behind distribution. A post may perform poorly because of source identity, copied text, low predicted engagement, topic fatigue, audience behavior, country restrictions, or timing. The dashboard may not distinguish among these causes. Even if the user conducts comparisons, the conclusions remain probabilistic. A pattern can suggest containment but rarely prove the exact mechanism.

This creates a structural asymmetry. The platform possesses the internal data and the user possesses only external symptoms. The platform can say, if challenged, that ranking systems are complex, that users did not engage, that no rule was violated, or that distribution varies naturally. The user cannot easily disprove these claims because the necessary data is unavailable. The platform controls both the distribution and the evidence of distribution.

This asymmetry is magnified for politically sensitive content. A platform has strong incentives not to reveal too much about how it handles state-media material, sanctions-sensitive sources, foreign influence signals, or legal compliance by region. Detailed disclosure could help actors evade enforcement. It could expose the platform to political criticism. It could reveal proprietary ranking systems. It could create legal complications. These reasons may be understandable, but they leave the ordinary user in the dark.

The problem is not unique to Facebook. All major platforms face similar measurement opacity. Search engines do not fully disclose ranking suppression. Video platforms do not fully disclose recommendation limits. App stores do not fully disclose discoverability decisions. Payment

processors do not always explain risk scoring. Email providers may filter messages into spam without sender knowledge. The modern information environment is filled with invisible distribution decisions whose effects are measurable only indirectly.

For the user posting RT and PressTV articles, the practical approach is triangulation. The dashboard provides one signal. Real viewers provide another. VPN direct-link tests provide another. Shareability checks provide another. Comparison across post types provides another. If all signals point in the same direction, confidence increases. If signals conflict, the user must accept uncertainty. The goal is not absolute proof but better inference.

A useful framework is to classify outcomes into four categories. First, no detectable restriction: RT and PressTV posts perform like comparable ordinary posts and are visible to outside viewers. Second, form-based restriction: full copied articles underperform excerpts and commentary regardless of source. Third, source-based restriction: RT or PressTV posts underperform comparable posts from other sources. Fourth, geographic restriction: posts visible in the United States are unavailable or less visible in specific foreign jurisdictions. These categories may overlap, but they help organize evidence.

The platform's failure to report certain data should itself be part of the analysis. A dashboard that cannot show countries, non-follower reach, recommendation eligibility, demotion status, or local-law restrictions is not sufficient for auditing public visibility. The user should not treat the dashboard as a complete truth instrument. It is a managerial interface, not a transparency system.

This measurement gap has broader civic implications. Democratic debate increasingly depends on platforms whose distribution decisions are hidden from speakers and audiences alike. A person can publish without knowing whether publication produced public communication. A society can believe it has free speech while speech is quietly sorted into silos. The absence of measurement becomes part of the governance structure.

The central conclusion of this section is therefore methodological and political. Methodologically, the Facebook dashboard cannot answer the siloing question by itself. Politically, the inability to answer that question is itself a feature of platform power. The user can count the visible reactions. He cannot count the people who were never allowed to see the post. In the age of soft censorship, the most important metric may be the one the platform does not show.

11. A Taxonomy of Modern Internet Restriction

Modern internet restriction is not one thing. It is a family of techniques that operate at different layers of the information system. Some are legal. Some are technical. Some are commercial. Some are algorithmic. Some are social. Some are explicit and easy to document. Others are hidden, probabilistic, and difficult to distinguish from ordinary low reach. A serious analysis of online censorship must therefore move beyond the single question, “Is it blocked?” That question belongs to an older model of information control. The better question is, “What kind of restriction is operating, at what layer, against which object, for which audience, and with what visibility to the speaker?”

The case of RT and PressTV reposting on Facebook shows why a taxonomy is necessary. The same article may be affected by several restriction systems at once. It may come from a state-linked outlet. It may be subject to sanctions or broadcast restrictions in some jurisdictions. It may be copied in full, creating copyright risk. It may appear on a private platform with state-media policies. It may be visible in the United States but restricted in parts of Europe. It may remain on the author’s page while being excluded from recommendation. It may be seen by followers but not by non-followers. It may be technically accessible by direct link but absent from normal feeds. None of these outcomes can be understood by the word “blocked” alone.

This taxonomy is not only descriptive. It is diagnostic. Each restriction type leaves different evidence. Hard blocking may produce a failed page load. Legal prohibition may produce an order, statute, sanctions notice, or regulator action. Platform removal may produce a notice. Geoblocking may be detected by foreign viewers. Demotion may appear only as reduced reach. Social containment may appear as repetitive engagement from a familiar audience. User unawareness may appear as the absence of any clear signal. To diagnose the system, one must know what traces each layer leaves.

The taxonomy also prevents exaggeration. Not every low-reach post is censored. Not every broadcast ban is a web block. Not every sanctions action prevents private quotation. Not every platform removal is government censorship. Not every copied article is protected commentary. Clear categories make the analysis stronger because they allow distinctions that the language of censorship often blurs.

The following categories describe the main forms of modern internet restriction relevant to the RT, PressTV, and Facebook case: hard blocking, legal prohibition, platform removal, geoblocking, demotion and distribution reduction, social containment, and user unawareness as a structural feature. These forms can operate separately, but their greatest power appears when they overlap.

11.1. Hard Blocking

Hard blocking is the clearest and oldest form of internet restriction. A user tries to reach a source and cannot. The domain fails to load. The app is unavailable. The video is blocked. The browser shows an error. The internet service provider refuses access. The government, court, network operator, platform, or institution has created a barrier that prevents ordinary access to the material.

Hard blocking is legible because it produces an encounter. The user experiences denial. The denial may be explained or unexplained, but the failure itself is obvious. A blocked page can be screenshotted. A blocked domain can be tested across networks. A blocked app can be identified in an app store. A seized domain may display a government notice. The restriction creates evidence.

The United States generally does not use hard blocking as the normal method for political news sources. A U.S. user can usually access foreign media outlets directly through the open web unless a specific legal, technical, or institutional layer intervenes. This is why a user in Kansas can often read RT, PressTV, or other foreign outlets without seeing a national firewall. Hard blocking is not the baseline American model for lawful political information.

But hard blocking does appear in narrower settings. Criminal domains may be seized. Child exploitation material may be blocked or removed. Malware domains may be filtered. Schools, libraries, workplaces, and public networks may block categories of content. App stores may remove apps. Platforms may block official pages or accounts. A country may order internet service providers to restrict a domain. These are all forms of hard blocking, but they differ in authority, scope, and purpose.

In the RT and PressTV case, hard blocking is only part of the story. RT may be blocked or unavailable through certain official channels, apps, platforms, or broadcast systems. PressTV may lose domain access, satellite access, platform accounts, or financial service access. Yet their content may still circulate through copied text, screenshots, mirrors, summaries, quotes, or reposts. Hard blocking the source does not necessarily hard block every derivative copy.

This is why hard blocking is insufficient as a diagnostic category. It tells us whether one access path is closed. It does not tell us whether the content survives in another form. It does not reveal whether copied content is demoted. It does not explain social siloing. It does not measure recommendation suppression. The blocked website is the visible wall. The modern system also contains invisible ceilings.

Hard blocking remains important because it is the most contestable form. It produces the clearest evidence and the strongest political reaction. But precisely because it is visible,

platforms and governments may prefer softer methods when possible. Hard blocking can create backlash. Containment can avoid it.

11.2. Legal Prohibition

Legal prohibition occurs when law declares that certain distribution, access, transaction, service, broadcast, or institutional relationship is unlawful. It may be created by statute, regulation, sanctions order, court order, broadcast authority, emergency decree, or administrative action. Unlike hard blocking, legal prohibition does not always produce immediate technical inaccessibility. A thing can be legally prohibited yet technically reachable. The law creates liability; the platform or service provider decides how to implement compliance.

Legal prohibition must be analyzed carefully because the prohibited object may vary. A government may prohibit broadcasting a channel, but not private discussion of that channel. It may sanction an entity, but not ban every completed informational material associated with it. It may prohibit providing services to an outlet, but not reading a newspaper article. It may restrict official distribution, but not scholarly quotation. The details matter.

RT illustrates a broad form of legal restriction in Europe after the Ukraine war. The concern was not merely one false statement but the distribution of Russian state-media channels as instruments of hostile information operations. This moved the object of restriction from individual content to institutional availability. Under such a framework, platforms may treat RT content as legally sensitive even when reposted by private users.

PressTV illustrates a different legal pattern. Its restrictions are often tied to sanctions, domain seizures, broadcast decisions, satellite carriage, and state-media designation. These mechanisms can limit official distribution and institutional support without necessarily creating a simple rule that no private citizen may quote or discuss PressTV content. The legal object is frequently the entity, transaction, service, or broadcast channel, not every reader's encounter with every text.

Legal prohibition therefore creates uncertainty for platforms. A platform operating globally must decide how to avoid liability in jurisdictions with different rules. It may implement a narrow restriction, such as removing official accounts. It may implement a broader restriction, such as blocking links. It may implement a country-specific restriction. It may demote content associated with the prohibited entity. It may remove content preemptively to avoid legal risk. The user may see only the final result.

For the private user, legal prohibition matters even when he is not directly targeted. A U.S.-based page may be lawful from the author's perspective, but Facebook must consider where the post is visible. If European law restricts availability of RT content to European audiences,

Meta may restrict or demote the post for those viewers. The author's domestic legal position does not determine the platform's global compliance burden.

Legal prohibition also interacts with intent. A user who quotes a restricted source to criticize it is not in the same position as someone who systematically republishes it as a substitute feed. But automated systems may not reliably infer intent. Platforms may rely on form and pattern: full-text copying, frequency, source concentration, lack of commentary, and audience behavior. The legal category becomes operationalized through technical proxies.

The core lesson is that legal prohibition is not identical to visible censorship. It may create no direct warning to the author. It may produce platform restrictions only in certain countries. It may affect official distribution more than private speech. It may be implemented inconsistently. Yet it changes the risk environment in which the post travels. Law does not need to block the author's browser to shape the post's audience.

11.3. Platform Removal

Platform removal occurs when Facebook, YouTube, X, TikTok, a search engine, an app store, or another private intermediary deletes, disables, hides, suspends, or terminates content or accounts under its own rules or under legal pressure. Removal is not always state censorship. It may be private policy enforcement. It may be copyright compliance. It may be sanctions compliance. It may be spam control. It may be political risk management. It may be a response to government requests. The source of pressure matters, but the user experiences removal as disappearance.

Removal is more visible than demotion. The post is gone. The account is suspended. The page is disabled. The video is unavailable. The app cannot be downloaded. The user may receive a notice, appeal option, or vague policy explanation. Even when the explanation is poor, removal creates an event. It can be documented, protested, or compared with other cases.

For RT, platform removal has been a major restriction layer. Official accounts, channels, apps, and pages associated with Russian state media have faced bans, regional restrictions, labeling, and removal from major platforms. These actions often apply to the entity rather than to every private user who discusses the entity. But once official distribution is removed, secondary circulation becomes more sensitive. If users reproduce official RT content in full, platforms may interpret that as an attempt to preserve or recreate a removed distribution channel.

For PressTV, platform removal has also appeared through account restrictions, domain disruptions, and state-media enforcement. The pattern may differ from RT, but the mechanism is similar: the platform restricts the official source or affiliated network. The content then migrates into unofficial forms, where enforcement becomes more ambiguous.

Platform removal can occur for copied articles through several routes. The platform may identify the copied text as belonging to a banned or restricted source. It may receive a copyright complaint. It may classify the post as spam or low-quality mass copying. It may receive a local-law request. It may find that the article contains separate policy violations. It may interpret the page as engaged in coordinated or inauthentic amplification. Each route may produce the same result: removal.

But the absence of removal does not prove absence of restriction. This is the central point of the paper. A post may remain online because the platform prefers not to remove it, lacks confidence, has not detected it, has not received a complaint, or treats it as commentary. It may still be demoted or siloed. Removal is only the visible end of the enforcement spectrum.

Platform removal also changes user behavior. Some users stop posting. Others try to evade detection. Others remove URLs, alter text, use screenshots, or move to private groups. These adaptations can make enforcement harder but may also make the user's behavior appear more suspicious. A platform may distinguish between transparent quotation and deliberate circumvention. Removing the URL is not automatically safer if the pattern looks like laundering restricted content.

The civic problem with platform removal is accountability. Major platforms function as essential distribution systems, but their decisions are governed by private terms, automated systems, opaque review, and business incentives. When they remove official state-media channels, they may be acting responsibly against propaganda or acting as private censors under political pressure. Both interpretations can coexist depending on the case. The taxonomy does not decide legitimacy automatically. It identifies the mechanism.

11.4. Geoblocking

Geoblocking restricts content by location. A post, website, video, app, or service may be visible in one country and unavailable in another. The user's location may be inferred from IP address, account settings, billing information, SIM card, GPS, device region, language, or other signals. Geoblocking is the technical expression of legal geography.

Geoblocking can be explicit. A viewer may see a notice that content is not available in their country. A video may display a regional restriction message. A post may show an unavailable-content notice. An app may not appear in a country's app store. These visible forms are easier to detect.

But geoblocking can also be subtle. The post may not appear in search. It may be absent from the page timeline for certain viewers. It may be hidden only from logged-in accounts in a region. It may open by direct link but not appear in recommendations. It may be blocked for some account types but not others. The viewer may not know that a missing post exists.

RT is a strong candidate for geoblocking because regional legal restrictions differ sharply. A U.S. author may see an RT-derived post. A European viewer may not. A Canadian viewer may see it differently. A UK viewer may be affected by platform policy rather than general web prohibition. The same post has several regional states.

PressTV may be geographically restricted through a different matrix. Its visibility may depend on sanctions regimes, platform policies toward Iranian state media, domain actions, local complaints, and country-specific platform decisions. It may not follow the same pattern as RT. Testing must therefore distinguish the two sources rather than assuming identical treatment.

A VPN can test crude geoblocking, but only imperfectly. It changes apparent IP location but not the full account identity. A real viewer in the target country is a better test. The strongest evidence comes from comparing several viewers in several jurisdictions and recording whether they can see the post by direct link, by visiting the page, through search, and through normal feed distribution.

Geoblocking is politically significant because it fragments the public sphere. A post may be part of public discourse in one country and absent in another. International debate becomes asymmetrical. The author may believe he is addressing foreign audiences who are legally or technically unable to receive the message. Readers in one country may not know that readers elsewhere see a different internet.

Geoblocking also allows platforms to avoid global deletion. Instead of removing content everywhere, they can restrict it where risk is highest. This may be a reasonable compromise in a world of conflicting laws. But it also makes censorship harder to see. The author outside the restricted region may experience no loss. The affected audience may not know what is missing. The platform alone sees the full map.

11.5. Demotion and Distribution Reduction

Demotion is the reduction of reach without removal. It is one of the most important forms of modern restriction because it is powerful, flexible, and hard to prove. The post stays online. The user can see it. Followers may see it. But the platform shows it to fewer people, excludes it from recommendations, ranks it lower in feeds, limits non-follower exposure, suppresses it in search, or makes it ineligible for certain surfaces.

Demotion can be justified in many ways. Platforms demote spam, clickbait, low-quality content, duplicate content, misinformation, borderline policy violations, engagement bait, content users report negatively, political content users may not want, or material from sources the platform considers risky. Some demotion is ordinary curation. A feed cannot show everything. Ranking requires preference. The controversy begins when demotion targets politically sensitive

material, disfavored sources, state-media content, or controversial viewpoints without adequate transparency.

For RT and PressTV reposting, demotion may occur for several reasons at once. The post may be political. It may be copied in full. It may be low-originality publisher text rather than original user commentary. It may come from a state-linked source. It may be associated with restricted entities. It may attract negative feedback from some users. It may be deemed unsuitable for recommendation. The platform may not need to remove it to limit its effect.

Demotion is difficult to detect because it looks like ordinary underperformance. A post with low reach may be demoted, or it may simply be uninteresting. A post with few shares may be suppressed, or readers may be reluctant to share it. A page with no growth may be contained, or the audience may be naturally small. The evidence is comparative: similar posts, posted under similar conditions, perform differently depending on source, format, or topic.

Distribution reduction can also be account-level. A page that repeatedly posts copied state-media content may be assigned lower recommendation eligibility. Future posts may start from a weaker position. The user may not receive a notice because the account is not banned. The page simply stops expanding. Existing followers may continue to see posts, creating the impression of normal operation.

Demotion is the core instrument of soft censorship because it controls public impact while preserving formal speech. It is difficult to challenge because the user does not know the counterfactual. It is difficult to document because the missing audience leaves no trace. It is difficult to regulate because platforms can describe it as ranking rather than censorship. Yet for public discourse, demotion can matter more than removal. A post that is allowed but never carried may be politically dead while formally alive.

The most honest conclusion is that demotion cannot usually be proven from one user's dashboard. It can only be inferred through patterns, external checks, and comparison. But that uncertainty should not make demotion invisible in theory. It is one of the central powers of platform governance.

11.6. Social Containment

Social containment occurs when speech circulates within a limited and self-reinforcing audience. The restriction may not be imposed directly by law or policy. It may emerge from social selection, algorithmic learning, audience behavior, and platform ranking. The same people see the posts. The same people react. The same interpretive community confirms the material. The post remains active but does not cross into broader publics.

Social containment is not always censorship. Communities naturally form. People follow pages that interest them. Foreign-policy posts will attract fewer readers than family photos. Long copied articles may attract fewer readers than short personal commentary. A small audience can be an ordinary result of niche content. But social containment becomes part of soft censorship when platform systems reinforce the boundary and prevent outward movement without making that boundary visible.

In the RT and PressTV case, social containment may explain why daily posting produces no visible problem. The posts may be reaching people already inclined to read such material. Those readers do not report it. They may like or comment. Their engagement teaches the platform that this content belongs in that audience. The same signal may also discourage distribution to users outside the cluster who are less likely to engage or more likely to object. The post remains alive inside the silo.

Social containment is psychologically deceptive because it provides feedback. A hard block produces silence. Social containment produces conversation, but only within the room. The author may feel public because he receives responses. Yet the responses may come from a narrow circle. The audience becomes a mirror. It confirms that the post exists while hiding the absence of outsiders.

This form of containment can affect all political groups. Mainstream liberals, conservatives, dissidents, antiwar communities, foreign-policy skeptics, religious groups, and activist networks can each inhabit silos. The issue is not that one side is uniquely enclosed. The issue is that platform systems organize attention around predicted engagement, and predicted engagement often means ideological familiarity. A controversial post may be shown to those who already want it and withheld from those who might be challenged by it.

Social containment also protects the platform. Content inside a small sympathetic audience produces less controversy than content exposed broadly. The platform can allow fringe, adversarial, or controversial material to exist while preventing it from entering mass circulation. This is a politically efficient compromise. It reduces accusations of direct censorship while limiting public impact.

The only way to test social containment is to measure outward movement. Does the post reach non-followers? Do unfamiliar names comment? Do shares carry it beyond the original circle? Can foreign viewers find it without a direct link? Do posts from restricted sources perform worse than original commentary? Does the audience grow, or does the same group repeatedly interact? Without these questions, engagement can be mistaken for public reach.

Social containment may be the most common fate of controversial political speech on large platforms. The speaker is not silenced. The speaker is categorized. The post is not erased. It is routed to the people most likely to keep it harmlessly inside the silo.

11.7. User Unawareness as a Structural Feature

User unawareness is not merely a side effect of modern restriction. It is a structural feature. The system works partly because the user cannot see enough to know what is happening. The author sees the post but not the distribution map. He sees reactions but not missing viewers. He sees aggregate reach but not always countries. He sees comments but not recommendation eligibility. He sees survival but not demotion. He sees continuity but not counterfactual reach.

This unawareness arises from several sources. Platform interfaces are simplified. Ranking systems are proprietary. Legal compliance may be confidential. Local-law restrictions may be disclosed only in limited ways. Demotion may be treated as an internal ranking decision. Privacy thresholds may hide geographic data. Account-level trust scores are not exposed. Automated classification is complex and often unexplained. The user receives an operational dashboard, not a governance audit.

User unawareness benefits platforms because it reduces conflict. If every demotion, source classification, recommendation exclusion, and regional restriction generated a notice, users would challenge decisions constantly. Platforms would face political controversy, legal questions, and pressure to explain their systems. Silence is easier. The platform can manage distribution without turning every ranking decision into a public dispute.

User unawareness also benefits governments indirectly. A state or regional authority can shape platform behavior through laws, sanctions, public pressure, or regulatory threat without every affected user seeing a censorship notice. The restriction becomes part of corporate compliance. The user experiences only a vague decline in reach or a missing foreign audience. Responsibility is dispersed.

For the user, unawareness creates unstable interpretation. He may conclude that because he has had no problem, there is no restriction. Or he may conclude that because reach is limited, censorship must be occurring. Both conclusions exceed the evidence. The system denies the data needed for certainty. This is why the correct response is structured skepticism: assume neither freedom nor censorship is proven by the interface alone.

User unawareness is especially important in the copied-article case. The author may believe removing URLs avoids restrictions. He may not know whether text recognition identifies the source. He may not know whether full reposting reduces distribution. He may not know whether EU viewers can see RT-derived posts. He may not know whether PressTV posts are

treated differently. He may not know whether non-followers receive the posts. His daily experience supplies only partial evidence.

This does not mean the platform is necessarily acting maliciously. Some opacity protects privacy, security, anti-spam systems, and enforcement integrity. Full transparency could help bad actors evade detection. But the same opacity also prevents legitimate users from understanding how public their speech really is. The system must balance safety against democratic transparency, and today that balance often favors platform control.

The final category in this taxonomy is therefore not a restriction technique in the narrow sense, but a condition that makes the other techniques hard to see. Hard blocking announces itself. Legal prohibition may be discoverable. Platform removal creates an event. Geoblocking may be tested. Demotion and social containment are harder. User unawareness is what allows them to remain ambiguous.

In the age of platform governance, censorship does not always require silence. It may require only that the speaker never learn the size of the room.

12. Conclusion: The Future of Censorship Is Partial Visibility

The central lesson of this paper is that censorship can no longer be understood only as removal, prohibition, or blocked access. Those older forms remain real, but they are no longer sufficient to describe how information moves through the contemporary digital public sphere. A post can survive and still be restricted. A source can be readable and still be distribution-limited. A speaker can remain active and still be enclosed. The future of censorship is not always the disappearance of speech. It is often the fragmentation of visibility.

This shift changes the meaning of freedom online. In the older model, the question was whether the state allowed a citizen to read, write, publish, or broadcast. In the platform model, the question becomes whether the post is carried, recommended, searchable, shareable, regionally available, and visible beyond an already-selected audience. Publication is no longer the same as publicness. Access is no longer the same as reach. A Facebook page may function as an archive, a room, a channel, or a cage depending on how the platform routes its contents.

The case of RT and PressTV shows this clearly because the sources occupy sensitive positions in the Western information order. They are foreign state-linked outlets associated with geopolitical conflict, sanctions, propaganda accusations, broadcast restrictions, domain actions, and platform enforcement. A U.S.-based user may be able to read them, copy them, and post their contents without encountering a direct federal block. But once those contents are placed on Facebook, they enter a global platform governed by many layers: U.S. law, European restrictions,

sanctions compliance, copyright law, Meta policy, local-law requests, ranking systems, account history, and audience behavior.

The result is not simple freedom or simple censorship. It is conditional circulation. An RT-derived post may appear in Kansas but fail to reach Europe. A PressTV-derived post may remain visible but be treated as state-linked or sanctions-sensitive. A full copied article may be allowed but demoted as low-originality material. A post may be visible by direct link but absent from feeds. A loyal audience may react while non-followers never see it. The post exists, but its public is uncertain.

This uncertainty is the defining political condition of platform speech. The user can no longer infer audience from publication. He must test, compare, verify, and measure. Even then, he may obtain only partial knowledge. The platform sees the full distribution map. The user sees shadows on the dashboard.

12.1. Why the Old Question "Is It Blocked?" Is No Longer Enough

The question "Is it blocked?" belongs to a simpler internet. It assumes that restriction is binary: either access is allowed or access is denied. That question still matters when a government blocks a domain, a platform removes an account, a regulator revokes a license, or a court orders a takedown. But it fails when restriction takes the form of demotion, geoblocking, regional unavailability, recommendation exclusion, duplicate-content suppression, source labeling, or social containment.

A modern post can have many states at once. It can be available in the United States and unavailable in the European Union. It can be visible to the author but not to non-followers. It can be readable through a direct link but not discoverable in search. It can be present on a page but excluded from recommendation. It can be legal to discuss but risky to monetize. It can be protected as commentary but vulnerable as full-text republication. It can be tolerated when small and restricted when it begins to spread.

Therefore the older question must be replaced by a sequence of better questions. Is the source blocked as a website, broadcaster, app, account, payment recipient, advertising entity, or search result? Is the content restricted globally or only in certain countries? Is the post removed, demoted, labeled, hidden, or simply not recommended? Can followers see it? Can non-followers see it? Can foreign viewers see it? Can it be shared? Does the dashboard report the relevant data? Does the author's view match the audience's view?

These questions reveal that "blocked" is often too crude. RT may be blocked or restricted as an official broadcaster in one jurisdiction, removed from a platform in another, demoted as a link, copied as text, and visible to a private audience. PressTV may be sanctioned, deplatformed,

disrupted, or labeled while still appearing in excerpts, archives, and reposts. The content moves through cracks between formal categories. It is neither fully available nor fully absent.

The old question also fails because it privileges access over circulation. In a platform environment, being accessible is not enough. A post can be accessible to anyone who already has the link while being practically invisible to everyone else. This is a profound change. Public communication once depended heavily on publication. Now it depends on distribution. Platforms control distribution more than publication. They may let the user speak while withholding the megaphone.

The more accurate question is therefore not “Is it blocked?” but “What is the post’s effective public?” A post with a public setting may still have a private effective audience. A post may be formally public but practically enclosed. A post may satisfy the author’s desire to publish while failing to enter public debate. The public is no longer a status selected by the user. It is an outcome produced by the platform.

This does not mean every low-reach post is censored. It means that the old binary language is no longer adequate. We need a vocabulary of layers: hard blocking, legal prohibition, platform removal, geoblocking, demotion, recommendation exclusion, social containment, copyright restriction, sanctions compliance, and user unawareness. Only with such a vocabulary can the modern information environment be described honestly.

12.2. The Rise of Jurisdictional and Algorithmic Fragmentation

The internet once promised a unified global public sphere. That promise has weakened. The contemporary internet is increasingly fragmented by law, platform policy, language, geography, sanctions, app stores, payment systems, recommendation algorithms, copyright systems, and national security concerns. Users may still see one interface, but behind that interface the world is divided into many publics.

Jurisdictional fragmentation means that law follows the viewer as much as the speaker. A post made in the United States may be judged by European restrictions when viewed in Europe, by British rules when viewed in the United Kingdom, by Canadian distribution norms when viewed in Canada, and by Australian sanctions policy when viewed in Australia. The platform must reconcile these conflicting legal environments. It does so by localizing visibility, removing content, demoting risk, or applying global policies that simplify compliance.

Algorithmic fragmentation means that even within one jurisdiction, users do not see the same public sphere. The platform ranks posts differently for different users. It predicts interest, risk, engagement, political sensitivity, source credibility, and likelihood of negative feedback. It shows one user a post and withholds it from another. It learns from social clusters and reinforces

them. The same nominally public page can appear differently to followers, non-followers, critics, foreign viewers, and logged-out visitors.

Together, jurisdictional and algorithmic fragmentation create a new condition: the individualized border. The border is no longer only between countries. It is also between accounts, audiences, feeds, trust categories, and recommendation systems. A post crosses one border but not another. The author cannot see most of these borders because they are not displayed as borders. They are experienced as ordinary variation in reach.

RT and PressTV content is especially vulnerable to this fragmentation because it triggers multiple classification systems. It is foreign. It is state-linked. It is politically charged. It may be sanctioned or restricted in some jurisdictions. It may be copied in full. It may be associated with narratives that platforms treat as civic, war-related, or foreign-influence content. A platform does not need one master rule to contain such material. Multiple small rules can produce the same outcome.

This is the architecture of partial visibility. No single actor needs to ban the post everywhere. One jurisdiction restricts the source. One platform removes official accounts. One algorithm demotes duplicate content. One audience cluster limits engagement. One dashboard hides geography. One sanctions regime makes companies cautious. One copyright system discourages full copying. The combined effect is a post that survives but does not travel freely.

Fragmentation also weakens public accountability. When a post is removed by a government order, responsibility is visible. When a post fails to circulate because of overlapping law, policy, ranking, and audience behavior, responsibility is diffused. The state can blame the platform. The platform can blame user preferences. The algorithm can be described as neutral optimization. The user can be told that no rule was violated. The result is restriction without a clear author.

This is why the modern debate over censorship must include platform architecture. Legal freedom is necessary but insufficient. A person may have the legal right to speak while lacking the practical ability to reach an audience. A public sphere governed by opaque ranking systems may preserve formal liberty while altering political reality. Fragmentation is not always censorship, but it is always power.

12.3. Practical Implications for Researchers, Writers, and Citizens

The first practical implication is that writers and citizens should distinguish between publication and reach. Posting is not proof of distribution. Seeing one’s own post is not proof of public visibility. Likes from familiar readers are not proof of broad circulation. A dashboard number is not proof of geographic availability. Anyone working with controversial, foreign, state-linked, or restricted sources must assume that visibility is layered until tested.

The second implication is that form matters. Full copied articles are more vulnerable than excerpts with commentary. A full repost can look like republication, mirroring, source laundering, or low-originality duplication. A short excerpt with clear attribution and analysis looks more like criticism, scholarship, or civic commentary. The more the user's own thought dominates the post, the stronger the position. The source becomes evidence rather than the entire message.

The third implication is that source identity matters even without URLs. Removing a link may reduce domain-level detection, but it does not erase provenance. Text can be recognized. Headlines can be matched. Patterns can be inferred. Readers can report. Repeated behavior can classify the page. The absence of a URL may even be interpreted as an attempt to bypass source controls if the post otherwise reproduces restricted content in full.

The fourth implication is that measurement must be comparative. A user should not ask only whether one post can be seen. He should compare ordinary posts, original commentary, RT excerpts, RT full reposts, PressTV excerpts, PressTV full reposts, and full copied articles from nonrestricted sources. Differences across these categories can reveal whether the issue is source identity, copied format, political topic, geographic restriction, or general page performance.

The fifth implication is that human verification matters more than artificial testing. A VPN can test crude geoblocking, but it does not reproduce a normal foreign Facebook user. A newly created foreign account may create false signals. Real viewers in target countries are better. They can test direct-link visibility, page visibility, search visibility, shareability, and warning messages under normal conditions. This kind of verification does not reveal everything, but it is closer to actual audience experience.

The sixth implication is that dashboards should be treated as partial instruments. If the dashboard does not show countries, non-follower reach, recommendation eligibility, local-law restrictions, or demotion status, it cannot answer the siloing question. Its numbers may be accurate and still incomplete. The most important missing metric may be the audience that never received the post.

The seventh implication is that researchers should treat platform opacity as part of the object of study. Missing data is not merely an inconvenience. It is a structural feature of platform power. Platforms govern public communication while disclosing only limited evidence about how that governance works. Any serious study of censorship, propaganda, state media, or digital public life must account for this asymmetry.

The eighth implication is ethical. Citizens should be able to read, quote, criticize, and analyze hostile or disfavored media. A democratic society cannot understand propaganda without

examining it. But there is a difference between analysis and uncritical relay. Reposting full articles daily from restricted state-linked outlets risks turning a citizen page into a substitute distribution channel. The stronger civic posture is transparent critical mediation: identify the source, quote what matters, compare claims, and add independent judgment.

The final practical implication is strategic. If the goal is to maximize durable visibility while reducing risk, the best practice is not to hide the source or mirror the article. It is to transform the material into original analysis. The user should make the post unmistakably his own speech about the source rather than the source's speech using his page as a vehicle. In the age of platform siloing, originality is not only a matter of style. It is a defense against containment.

12.4. Final Claim: The Post Exists, but the Public May Not

The final claim of this paper is simple: the existence of a post no longer guarantees the existence of a public. A post can be present without being public in any meaningful sense. It can sit on a page, visible to the author and a familiar circle, while being absent from foreign jurisdictions, non-follower feeds, search surfaces, recommendation systems, and wider debate. The post exists, but the public may not.

This claim does not require paranoia. It does not assume that every platform decision is malicious or every low-reach post is censored. It requires only realism about how modern platforms work. Facebook and similar systems do not merely display speech. They distribute speech according to rules, predictions, incentives, risks, and legal demands. Distribution is the new terrain of censorship. The border is no longer only at the website address. It is inside the feed.

For the user posting RT and PressTV articles, the most defensible conclusion is neither complacency nor certainty of persecution. The evidence shows that daily posting may continue without overt punishment. It does not show that the posts are globally visible. It does not show that foreign viewers can see them. It does not show that non-followers receive them. It does not show that copied full articles are treated like original commentary. It does not show that future enforcement will match past tolerance. The correct conclusion is disciplined uncertainty.

Disciplined uncertainty leads to better practice. Test visibility. Use real foreign viewers. Compare post categories. Track non-follower reach where possible. Avoid full-text mirroring when commentary will do. Preserve attribution. Add analysis. Treat the dashboard as partial. Recognize that a post can survive and still be contained. Refuse both naive trust and unmeasured suspicion.

The broader political conclusion is more troubling. A society can retain formal free speech while losing transparent public reach. Citizens may still speak, but platforms decide how far speech travels. Governments may avoid direct censorship while shaping corporate compliance.

Audiences may be fragmented into legal, algorithmic, and social silos. Writers may mistake publication for communication. The public sphere may remain visually intact while internally partitioned.

This is the new geography of internet visibility. The old map showed open and closed countries, blocked and unblocked sites, banned and permitted sources. The new map is finer and stranger. It shows partial publics, conditional posts, localized absences, demoted speech, missing metrics, and audiences that exist only for some viewers. The future of censorship is not always silence. It is partial audibility.

The post exists. The author sees it. The archive records it. The familiar audience may respond. But the decisive democratic question remains unanswered until tested: who else can see it, and who has been quietly removed from the room?

13. References

13.1. Legal and Regulatory Sources

1. Freedom House. United States: Freedom on the Net 2025. Freedom House, 2025.
2. Federal Communications Commission. Children's Internet Protection Act. FCC Consumer Guide, updated 2024.
3. U.S. Department of Justice. Child Sexual Abuse Material. DOJ, 2023.
4. 18 U.S.C. § 2252. Certain activities relating to material involving the sexual exploitation of minors. Legal Information Institute, Cornell Law School.
5. American Library Association. First Amendment and Censorship. ALA Office for Intellectual Freedom.
6. 31 C.F.R. § 560.210. Exempt transactions. Iranian Transactions and Sanctions Regulations. Legal Information Institute, Cornell Law School.
7. 31 C.F.R. § 560.315. Information or informational materials. Electronic Code of Federal Regulations.
8. Council of the European Union. EU imposes sanctions on state-owned outlets RT/Russia Today and Sputnik's broadcasting in the EU. Press release, March 2, 2022.
9. Council Regulation (EU) No 833/2014. Restrictive measures in view of Russia's actions destabilising the situation in Ukraine. EUR-Lex consolidated legal text.
10. Council Regulation (EU) 2022/350. Amendment to Regulation (EU) No 833/2014 concerning restrictive measures in view of Russia's actions destabilising the situation in Ukraine. Official Journal of the European Union, 2022.
11. Court of Justice of the European Union. Case C-67/25, Traugott Ickeroth; R, N and K v Staatsanwaltschaft Saarbrücken. Judgment concerning the scope of the RT broadcasting prohibition under Article 2f(1) of Regulation No 833/2014, July 2, 2026.
12. Ofcom. Ofcom revokes RT's broadcast licence. March 18, 2022.
13. Canadian Radio-television and Telecommunications Commission. Broadcasting Decision CRTC 2022-68: Russia Today and RT France removed from the list of non-Canadian programming services authorized for distribution. March 16, 2022.
14. Government of Canada. RT and RT France can no longer be distributed by Canadian television service providers. CRTC news release, March 16, 2022.

15. Reuters. Moldova blocks access to more than 20 Russian media websites. October 24, 2023.

16. Radio Free Europe/Radio Liberty. Moldova Blocks Access To More Than 20 Russian Media Websites. October 24, 2023.

17. Ofcom. Revocation: Provider of the Service Press TV Limited. January 20, 2012.

18. Reuters. Britain pulls the plug on Iran's Press TV. January 20, 2012.

19. Council of the European Union. Iran: EU adopts additional sanctions against perpetrators of serious human rights violations. November 14, 2022.

20. European Union Sanctions Tracker. Press TV designation under the Iran sanctions regime; designation date November 14, 2022.

21. U.S. Department of the Treasury, Office of Foreign Assets Control. Treasury Sanctions Iranian Officials and Companies Connected to Continued Repression in Iran. September 15, 2023.

22. OFAC Sanctions List Search. Press TV entry, Specially Designated Nationals and Blocked Persons List.

23. Global Affairs Canada. Backgrounder: Additional sanctions against Iran. October 3, 2022.

24. Australian Minister for Foreign Affairs. Targeted sanctions in response to human rights violations in Iran. September 13, 2023.

25. U.S. Department of Justice. United States Seizes Websites Used by the Iranian Islamic Radio and Television Union and Kata'ib Hizballah. June 2021.

26. Al Jazeera. US seizes three dozen websites used for Iranian disinformation. June 23, 2021.

27. United States Institute of Peace, The Iran Primer. U.S. Seizes Websites Tied to Iran. July 7, 2021.

13.2. Platform Policy and Transparency Sources

28. Meta Transparency Center. Content Restrictions Based on Local Law. Meta Platforms.

29. Meta Transparency Center. How we assess reports of content violating local law. Meta Platforms.

30. Meta Transparency Center. Content Restrictions Case Studies. Meta Platforms.

31. Meta Transparency Center. Authentic Identity Representation. Community Standards, Meta Platforms.

32. Meta Transparency Center. Account Integrity. Community Standards, Meta Platforms.

33. Facebook Help Center. See Page Insights on Facebook. Meta Platforms.

34. Meta Business Help Center. Troubleshoot When Demographic Data Isn't Available in Meta Business Suite. Meta Platforms.

35. Meta Business Help Center. Insights on your Facebook Page. Meta Platforms.

36. Reuters. Meta bans Russian state media for foreign interference. September 17, 2024.

37. Radio Free Europe/Radio Liberty. Meta Bans RT, Other Russian State Media Networks. September 17, 2024.

38. UNESCO. Content Restrictions Based on Local Law. World Trends in Freedom of Expression and Media Development resource.

39. Article 19. Content Moderation and Freedom of Expression Handbook. 2023.

40. Reuters. EU preliminarily finds Meta, TikTok in breach of transparency obligations. October 24, 2025.

13.3. Journalism, Reporting, and Case Materials

41. The Guardian. Russia threatens further crackdown on British media after Ofcom bans RT. March 18, 2022.

42. Axios. Britain bans Russian state-backed media outlet RT. March 18, 2022.

43. The Guardian. BitChute to continue hosting Kremlin-backed RT channel. March 9, 2022.

44. Reuters. EU ambassadors agree to place four Russian media outlets on sanctions list. May 15, 2024.

45. Politico. RT running covert military supply operation for Russian forces, State Dept. says. September 13, 2024.

46. Reuters. Meta bans Russian state media for foreign interference. September 17, 2024.

47. The Guardian. Iran's Press TV loses UK licence. January 20, 2012.

48. Voice of America. Sanctioned for Airing Forced Confessions, Iran's Propaganda Mouthpiece Cries Press Freedom. September 18, 2023.

49. The Guardian. Australia to impose sanctions on Iranian state media over broadcast of forced confessions. September 12, 2023.

50. ABC News Australia. Sharma says sanctions warning to university over Iran state broadcaster PressTV. June 27, 2025.

51. The Guardian. Meta and Snapchat blocking Saudi dissidents' accounts. May 22, 2026.

13.4. Academic Literature on Platform Governance, Moderation, and Networked Visibility

52. Gillespie, Tarleton. Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media. Yale University Press, 2018.

53. Roberts, Sarah T. Behind the Screen: Content Moderation in the Shadows of Social Media. Yale University Press, 2019.

54. Suzor, Nicolas. Lawless: The Secret Rules That Govern Our Digital Lives. Cambridge University Press, 2019.

55. DeNardis, Laura. The Internet in Everything: Freedom and Security in a World with No Off Switch. Yale University Press, 2020.

56. Napoli, Philip M. Social Media and the Public Interest: Media Regulation in the Disinformation Age. Columbia University Press, 2019.

57. Benkler, Yochai, Robert Faris, and Hal Roberts. Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics. Oxford University Press, 2018.

58. Zuboff, Shoshana. The Age of Surveillance Capitalism. PublicAffairs, 2019.

59. Gorwa, Robert. The Politics of Platform Regulation: How Governments Shape Online Content Moderation. Oxford University Press, 2024.

60. Douek, Evelyn. Governing Online Speech: From Posts-as-Trumps to Proportionality and Probability. Columbia Law Review, 2022.

61. Keller, Daphne. Amplification and Its Discontents. Knight First Amendment Institute, 2021.

62. Marwick, Alice, and Rebecca Lewis. Media Manipulation and Disinformation Online. Data & Society Research Institute, 2017.

63. Bradshaw, Samantha, and Philip N. Howard. The Global Disinformation Order: 2019 Global Inventory of Organised Social Media Manipulation. Oxford Internet Institute, 2019.

64. Starbird, Kate. Disinformation's Spread: Bots, Trolls and All of Us. Nature, 2019.

65. Phillips, Whitney, and Ryan M. Milner. You Are Here: A Field Guide for Navigating Polarized Speech, Conspiracy Theories, and Our Polluted Media Landscape. MIT Press, 2021.

66. Gillespie, Tarleton. Platforms Are Not Intermediaries. Georgetown Law Technology Review, 2018.

67. Gorwa, Robert, Reuben Binns, and Christian Katzenbach. Algorithmic Content Moderation: Technical and Political Challenges in the Automation of Platform Governance. Big Data & Society, 2020.

68. Helberger, Natali, Kari Karppinen, and Lucia D'Acunto. Exposure Diversity as a Design Principle for Recommender Systems. Information, Communication & Society, 2018.

69. Rieder, Bernhard, and Jeanette Hofmann. Towards Platform Observability. Internet Policy Review, 2020.

70. Gillespie, Tarleton. The Politics of Platforms. New Media & Society, 2010.

13.5. Studies and Reports on Foreign Influence, State Media, and Coordinated Information Operations

71. Kurek, Laura, Kevin Zheng, Eric Gilbert, and Ceren Budak. Outsourcing an Information Operation: A Complete Dataset of Tenet Media's Podcasts on Rumble. arXiv, 2025.

72. Cinus, Federico, Marco Minici, Luca Luceri, and Emilio Ferrara. Exposing Cross-Platform Coordinated Inauthentic Activity in the Run-Up to the 2024 U.S. Election. arXiv, 2024.

73. Burnat, Florian A. D., and Brittany I. Davidson. The Accountability Paradox: How Platform API Restrictions Undermine AI Transparency Mandates. arXiv, 2025.

74. Vériter, Sophie L. The Collective Securitization of Disinformation and the EU's Ban on Russia Today and Sputnik. International Affairs, 2025.

75. Lendvai, György F. Media in War: An Overview of the European Restrictions on Russian Media. European Papers, 2024.

76. Finelli, Francesco. Proportionate Sanctions? Balancing Effectiveness and Fundamental Rights Before the CJEU. European Papers, 2026.

77. United24 Media. EU Court Confirms Blanket Ban on All RT Russian State Media Content. July 2026.

78. EU News. Ban on broadcasting Russia Today: the EU Court clarifies that it also applies to free websites. July 2, 2026.

79. Global Sanctions. European Court of Justice interprets “operator” in Russia Regulation 833 broadcast prohibition. July 2026.

80. IEU Monitoring. EU Court confirms Russia Today broadcast ban covers free public websites. July 2026.

The author has published over 5,600 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.